FIVE ESSAYS IN MACROECONOMICS

INAUGURAL-DISSERTATION

zur Erlangung des akademischen Grades
eines Doktors der Wirtschaftswissenschaft
doctor rerum politicarum
(Dr. rer. pol.)

am Fachbereich Wirtschaftswissenschaft
der Freien Universität Berlin

vorgelegt von
Fabian Georg Ludwig Seyrich

Berlin, Juni 2024

Dekanin:	Prof. Dr. Natalia Kliewer
Erstgutachter:	Prof. Dr. Alexander Kriwoluzky Freie Universität Berlin
Zweitgutachter:	Prof. Dr. Christian Bayer Universität Bonn
Tag der Disputation:	09. Juli 2024

Acknowledgements

I would like to express my deepest gratitude to those who have supported me throughout my dissertation journey and without which this dissertation would have not been possible.

I am immensely thankful to my first supervisor, Alex. His supervision—a mixture between trust and laissez-faire while always being ready and available to give invaluable help and counsel when needed—was ideal for me. As a supervisor and boss, he struck the perfect balance between pushing me for academic excellence and creating a fun work environment by mixing sharp intellectual debates with witty comments—both of which has been highly educational for me. He also was a huge support during the stressful times before, during, and after the academic job market. I am also very grateful that he accepted me as a PhD student midway through my dissertation, despite already having a full roster of PhD students, as well as for the opportunity to do a research stay in Chicago and to go on countless conferences.

I also owe a special thanks to Christian, my second supervisor. Needless to say, how helpful and instructive it was to be able to work with the leading expert in the field of my thesis throughout my dissertation. Throughout working on our joint projects and by receiving plenty of counsel on my other projects, I learned a great deal about economic intuition, economic modelling, and coding. I am especially grateful for how generous he was with his time, sometimes going through codes line by line for hours or to help me to improve my job market talk.

I am also very grateful to Gernot, my unofficial third supervisor. Working with him was extremely educational and I especially enjoyed our close collaboration on the HANK^2 paper. Through this experience, I learned a lot not only about (international) macroeconomics but also about how to embed research questions into bigger contexts, and about the art of how to transform ideas into well-polished papers.

Many thanks also go to my other co-authors. I am extremely glad that after several failed trials, Oliver and I eventually figured out a topic to jointly work on after we had already planned to write a joint paper over beers at Café Vienna in 2019. Not only because the BHANK journey so far was a hell of a journey, but much more importantly, because it showed me that academic working does not at all need to feel like working but rather feels like spending time with a great friend.

Hannah deserves special mentions and thanks for sharing the burden of writing (and rewriting at least 871358 times) our first paper, getting the first journal rejection, getting the first revise & resubmit, and eventually publish the first paper. What an exciting journey this first paper was, and I am very happy to have been able to share it with you!

A heartfelt thanks to my office buddies who made the PhD journey as enjoyable as it was. Andrés, Jonathan, Jacopo, and again Oliver helped me to enjoy my first year in Mannheim—either by easing the pain of studying micro and metrics proofs in doing it together ("Can we do that on the board?") or going to Nelson. At the DIW, I was happy and lucky to meet Gökhan, Frederik, and Ben. The few office days without a joint mensa feast reminded me of how much tougher this journey would have been without friends in the office. May we go on many more Barcelona revivals together in the future and share the patatas bravas!

Of course, a big thank you also goes to the many people at the Macro Department of DIW Berlin, the DIW Graduate Center and the Berlin School of Economics who have helped me in many ways throughout the years.

Lastly, I am profoundly grateful to the people outside of academia who supported this journey. The greatest thanks of all goes to my parents, who have always supported me in all aspects of life including every step in my education. I could not have wished for any better parental support! Another big thanks go to my older brothers Sebastian, Jonathan, and Maximilan. Already in high school I could always ask them whenever I needed help with school (besides all the help they provided me in other aspects of life). Probably, they also deserve thanks for putting me off physics which helped me to end up at economics in the first place.

Finally, a big thank you to Ellen. I am immensely grateful for her personal support that made the PhD journey seem like a very easy journey. From joint lunch breaks in Steglitz, putting up with my hourless calls through thin walls while working room-to-room in Stuttgart, where I endlessly rambled on about RANK/HANK/TANK ("How is that even working?"), to working late-hour-shifts building-to-building with shared journeys home and setting off for work together in the morning: She always provides the perfect framework for my working days.

Thank you all for being part of this journey with me!

Berlin, June 2024
Fabian Seyrich

Declaration of Co-Authorship and Publications

This dissertation consists of five (working) papers. The contribution in conception, implementation, and drafting of the five chapters can be summarized as follows:

1. Climate HANK: Carbon Shocks, Fiscal Policy, and Double Dividends
   Contribution: 100%

2. Unconventional Fiscal Policy in a Heterogeneous-Agent New Keynesian Model
   with Hannah Seidl
   Contribution: 50%
   This version was published as
   Seidl, H. and F. Seyrich (2023). Unconventional Fiscal Policy in a Heterogeneous-Agent New Keynesian Model *Journal of Political Economy Macroeconomics* 1.4, 633–664.

3. A HANK^2 model of monetary unions
   with Christian Bayer, Alexander Kriwoluzky, and Gernot Müller
   Contribution: 25%
   This version was published as
   Bayer, C., A. Kriwoluzky, G. Müller, and F. Seyrich (2024). A HANK^2 model of monetary unions *Journal of Monetary Economics* 103579.

4. A Behavioral Heterogeneous Agent New Keynesian Model
   with Oliver Pfäuti
   Contribution: 50%

5. Bad Luck or Bad Decisions? Macroeconomic Implications of Persistent Heterogeneity in Cognitive Skills and Overconfidence
   with Oliver Pfäuti and Jonathan Zinman
   Contribution: 33.3%

Contents

List of Figures

List of Tables

Summary

Chapter 1: In this paper, I analyze the business cycle and distributional impact of carbon pricing policies. To this end, I develop a heterogeneous-agent New Keynesian model with a carbon market. I show that an adverse carbon shock, that is a temporary decrease in carbon certificates, increases inflation and reduces economic activity such as output, inflation and investment. I show that the economic costs of higher carbon prices are borne relatively evenly across the income and wealth distribution. With the additional revenue from carbon pricing, fiscal policy can mitigate the macroeconomic and welfare effects of higher carbon prices. If these revenues are immediately paid back to households in the form of lump-sum transfers, higher carbon prices can even have a welfare-enhancing effect that goes beyond the welfare gain from lower emissions—a double dividend result.

Chapter 2: In this paper, which is joint work with Hannah Seidl, we derive a perfect substitutability result between fiscal policy and monetary policy in a heterogeneous-agent New Keynesian model. We derive in closed form three simple conditions for consumption taxes, labor taxes, and the government debt level that are sufficient to induce the same consumption and labor supply of each household and, thus, the same allocation as interest rate policies. The intuition is that consumption taxes and labor taxes replicate the effects monetary policy has on the wedges in the first-order conditions of households. Combining these with debt policies can then in addition generate equivalence in the budget constraint of each household. Given our result, fiscal policy can replicate any allocation that hypothetically unconstrained monetary policy would generate when monetary policy is constrained by a binding lower bound, a monetary union, or an exchange rate peg.

Chapter 3: In this paper, which is joint work with Christian Bayer, Alexander Kriwoluzky, and Gernot Müller, we analyze how a monetary union alters the impact of business cycle shocks at the household level. To this end, we develop a Heterogeneous Agent New Keynesian model of two countries (HANK^2). We show in closed form that, first, when aggregated across both countries and households, macroeconomic dynamics are independent of whether there is a monetary union or not. Second, the overall effect of the shock for households in specific brackets of the wealth (and income) distribution, aggregated across countries of residence, does not depend on whether monetary union is in place or not. In other words, the monetary union itself does not shift the impact of the shock *vertically* across wealth classes. Instead, third, a monetary union shifts the impact of the shock *horizontally* across borders within the brackets of the wealth distribution. Calibrating the

model to the euro area reveals that a monetary union alters the impact of shocks most strongly in the tails of the wealth distribution but leaves the middle class almost unaffected.

Chapter 4: In this paper, which is joint work with Oliver Pfäuti, we develop a behavioral heterogeneous-agent New Keynesian model. The model features a New Keynesian core and we allow for household heterogeneity and incomplete markets as well as bounded rationality in the form of cognitive discounting. Under cognitive discounting, households' expectations underreact to macroeconomic news, consistent with what we find in survey data. Our model accounts for recent important empirical findings on the transmission mechanisms and effectiveness of monetary policy. In particular, monetary policy affects household consumption to a large extent through indirect effects which tend to amplify the effects of conventional monetary policy on consumption as the incomes of households that exhibit higher marginal propensities to consume are more exposed to aggregate income fluctuations induced by monetary policy. Announcements of future monetary policy changes, in contrast, have relatively weak effects on current economic activity, and the economy remains stable at the effective lower bound. In contrast to existing models, we account for these facts simultaneously without having to rely on a specific monetary or fiscal policy. Given its consistency with the empirical findings about the monetary transmission, we use our model to revisit the policy implications of inflationary supply shocks. We uncover a novel amplification channel of supply shocks: Cognitive discounting and the unequal exposure of households as well as their interaction amplify the inflationary effects of supply shocks such that inflation increases 2.5 times as strongly as in models abstracting from these features.

Chapter 5: In this paper, which is joint work with Oliver Pfäuti and Jonathan Zinman, we analyze how heterogeneity within behavioral biases explains heterogeneity in households' saving behavior which in turn has important implications for macroeconomic fluctuations and fiscal policies. Using survey data, we empirically document a systematic relationship between households' cognitive skills and their savings behavior: Households with lower cognitive skills are more likely to be overly optimistic with respect to their future financial situation, they are more likely to be hand-to-mouth, and they tend to be overconfident with respect to their cognitive skills. We then introduce heterogeneity in cognitive skills and overconfidence into a Heterogeneous Agent New Keynesian model and show that our model can account for our empirical findings. Overconfidence proves to be the key innovation, driving households to spend instead of precautionary save and producing empirically realistic wealth distributions and hand-to-mouth shares and MPCs across the income distribution. We show that accounting for heterogeneity in overconfidence has important implications for fiscal policy: Providing liquidity is less effective in bringing households away from the borrowing constraint, low-income benefits lead to less crowding out of self insurance, and the optimal level of public debt is substantially lower than in standard models.

Zusammenfassung

Kapitel 1: In diesem Artikel analysiere ich die Konjunktur- und Verteilungseffekte einer Kohlenstoffpreispolitik. Zu diesem Zweck entwickle ich ein neukeynesianisches Modell mit heterogenen Agenten und einem Markt für Kohlenstoffdioxid. Ich zeige, dass ein negativer Kohlenstoffdioxidschock, d.h. eine vorübergehende Verringerung der Kohlenstoffzertifikate, die Inflation erhöht und die Wirtschaftstätigkeit wie Produktion, Inflation und Investitionen verringert. Ich zeige, dass die wirtschaftlichen Kosten höherer Kohlenstoffpreise relativ gleichmäßig über die Einkommens- und Vermögensverteilung getragen werden. Mit den zusätzlichen Einnahmen aus der Bepreisung von Kohlenstoffemissionen kann die Fiskalpolitik die makroökonomischen und Wohlfahrtseffekte höherer Kohlenstoffpreise abmildern. Wenn diese Einnahmen sofort in Form eines Pauschaltransfers an die Haushalte zurückgezahlt werden, können höhere Kohlenstoffpreise sogar einen wohlfahrtssteigernden Effekt haben, der über den Wohlfahrtsgewinn durch niedrigere Emissionen hinausgeht—ein doppeltes-Dividenden-Ergebnis.

Kapitel 2: In dieser Arbeit, die in Zusammenarbeit mit Hannah Seidl entstanden ist, leiten wir ein Substituierbarkeitsergebnis zwischen Fiskal- und Geldpolitik in einem Neukeynesianischen Modell mit heterogenen Agenten her. Wir zeigen analytisch, dass drei einfache Bedingungen für Konsumsteuern, Arbeitssteuern und die Höhe der Staatsverschuldung ausreichen, um den gleichen Konsum und das gleiche Arbeitsangebot jedes Haushaltes und damit die gleiche Allokation wie durch eine Zinspolitik herbeizuführen. Die Intuition ist, dass Konsum- und Arbeitssteuern die Auswirkungen der Geldpolitik auf die Keile in den Bedingungen erster Ordnung der Haushalte nachbilden. Kombiniert man diese mit der angemessenen Schuldenpolitik, so kann man zusätzlich Äquivalenz in der Budgetbeschränkung jedes Haushalts erzeugen. Unsere Ergebnisse implizieren, dass die Fiskalpolitik jede Allokation nachbilden kann, die eine hypothetisch uneingeschränkte Geldpolitik erzeugen würde, wenn die Geldpolitik durch eine verbindliche Untergrenze, eine Währungsunion oder eine Wechselkursanbindung eingeschränkt ist.

Kapitel 3: In diesem Papier, das in Zusammenarbeit mit Christian Bayer, Alexander Kriwoluzky und Gernot Müller entstanden ist, analysieren wir, wie eine Währungsunion die Auswirkungen von Konjunkturschocks auf der Ebene der privaten Haushalte verändert. Zu diesem Zweck entwickeln wir ein neukeynesianisches Modell mit heterogenen Agenten und zwei Ländern (HANK^2). Wir zeigen analytisch, dass erstens die makroökonomische Dynamiken, wenn diese über beide Länder und Haushalte aggregiert werden, unabhängig

davon sind, ob es eine Währungsunion gibt oder nicht. Zweitens hängt der Gesamteffekt des Schocks für Haushalte entlang der Vermögens- (und Einkommens-) Verteilung, aggregiert über beide Länder, nicht davon ab, ob es eine Währungsunion gibt oder nicht. Mit anderen Worten, die Währungsunion verschiebt die Auswirkungen des Schocks nicht zwischen verschiedenen Vermögensklassen. Stattdessen, drittens, verschiebt eine Währungsunion die Auswirkungen des Schocks horizontal über die Grenzen innerhalb der Klammern der Vermögensverteilung. Eine Kalibrierung des Modells auf den Euroraum zeigt, dass eine Währungsunion die Auswirkungen von Schocks am stärksten an den Enden der Vermögensverteilung verändert, während die Mittelschicht nahezu unberührt bleibt.

Kapitel 4: In diesem Papier, das in Zusammenarbeit mit Oliver Pfäuti entstanden ist, entwickeln wir ein verhaltensökonomisches, neukeynesianisches Modell mit heterogenen Agenten. Das Modell hat einen neukeynesianischen Kern, und wir berücksichtigen Heterogenität der Haushalte und unvollständige Märkte sowie begrenzte Rationalität in Form von kognitiver Diskontierung. Bei kognitiver Diskontierung reagieren die Erwartungen der Haushalte unterdurchschnittlich auf makroökonomische Neuigkeiten, was mit Ergebnissen aus Umfragedaten übereinstimmt. Unser Modell erklärt die jüngsten wichtigen empirischen Erkenntnisse über Transmissionsmechanismen und Wirksamkeit der Geldpolitik. Insbesondere beeinflusst die Geldpolitik den Konsum der privaten Haushalte zu großen Teilen durch indirekte Effekte, die dazu neigen, die Auswirkungen der konventionellen Geldpolitik auf den privaten Konsum zu verstärken, da die Einkommen der Haushalte, die eine höhere marginale Konsumneigung aufweisen, stärker den durch die Geldpolitik verursachten Schwankungen des Gesamteinkommens ausgesetzt sind. Die Ankündigung künftiger geldpolitischer Politike hat dagegen nur relativ schwache Auswirkungen auf die aktuelle Wirtschaftstätigkeit, und die Ökonomie bleibt an der Nullzinsgrenze stabil. Im Gegensatz zu bestehenden Modellen berücksichtigt unser Modell diese Tatsachen gleichzeitig, ohne dass es eine bestimmte Geld- oder Fiskalpolitik dazu braucht. Angesichts der Konsistenz mit den empirischen Erkenntnissen über die Transmission von Geldpolitik verwenden wir unser Modell, um die Auswirkungen von inflationären Angebotsschocks zu analysieren. Dabei entdecken wir einen neuen Verstärkungskanal für Angebotsschocks: Kognitive Diskontierung und ungleiche Exposition der Haushalte sowie ihr Zusammenspiel verstärken die inflationären Auswirkungen von Angebotsschocks, so dass die Inflation 2,5 Mal so stark ansteigt wie in Modellen, die von diesen Merkmalen abstrahieren.

Kapitel 5: In dieser Arbeit, die in Zusammenarbeit mit Oliver Pfäuti und Jonathan Zinman entstanden ist, analysieren wir, wie die Heterogenität innerhalb der Verhaltensbeschränkungen die Heterogenität im Sparverhalten der Haushalte erklärt, was wiederum wichtige Auswirkungen auf makroökonomische Schwankungen und die Fiskalpolitik hat. Anhand von Umfragedaten dokumentieren wir empirisch eine systematische Beziehung zwischen den kognitiven Fähigkeiten von Haushalten und ihrem Sparverhalten: Haushalte mit geringeren kognitiven Fähigkeiten sind mit größerer Wahrscheinlichkeit übermäßig optimistisch in Bezug auf ihre zukünftige finanzielle Situation, sie leben eher von der Hand in den Mund und sind tendenziell übermäßig zuversichtlich in Bezug auf ihre kognitiven Fähigkeiten. Anschließend führen wir Heterogenität bezüglich kognitiven Fähigkeiten und bezüglich Selbstüberschätzung der Haushalte in ein Neukeynesianisches Modell mit heterogenen Agenten ein und zeigen, dass unser Modell unsere empirischen Ergebnisse erklären kann. Selbstüberschätzung erweist sich als die entscheidende Neuerung, die die Haushalte dazu veranlasst, Geld auszugeben, anstatt vorsorglich zu sparen, und die zu empirisch realistischen Vermögensverteilungen, Anzahl an Haushalten, die von der Hand in den Mund leben und Konsumneigungen entlang der Einkommensverteilung führt. Wir zeigen, dass die Berücksichtigung der Heterogenität der Selbstüberschätzung wichtige Auswirkungen auf die Fiskalpolitik hat: Die Bereitstellung von Liquidität ist weniger wirksam, um die Haushalte von der Kreditgrenze wegzubringen, Sozialleistungen bei niedrigem Einkommen führen zu einer geringeren Verdrängung der Selbstversicherung, und die optimale Höhe der Staatsverschuldung ist niedriger als in Standardmodellen.

Introduction and overview

New Keynesian (NK) models are the cornerstone of monetary economics, as they provide insights into the effects of monetary policy and its role in stabilizing the business cycle. To this end, there has long been a consensus that focusing on a representative agent and thus abstracting from household heterogeneity (and other forms of heterogeneity) is innocuous when it comes to understanding business cycle fluctuations and the effects and transmission mechanism of stabilization policies. However, with the Great Financial Crisis, academics and policymakers likewise began to be increasingly unsatisfied with the simplifying assumption of representative agents. First, because representative agent models do not capture the business cycle effects of certain shocks such as credit crunches or income uncertainty shocks which, however, likely played a role during the Great Financial Crisis. Second, because these models do not provide a satisfactory description of the transmission mechanism of monetary policy and at best offer limited scope for analyzing fiscal policy. And thirdly, because by their very nature, these models cannot speak to the distributional effects of business cycle shocks and stabilization policies and the repercussions of these distributional effects on the macroeconomy.

Heterogeneous-Agent New Keynesian (HANK) models address these limitations by providing a rich representation of income and wealth distributions, as well as more realistic consumption patterns, including heterogeneous and potentially large marginal propensities to consume (Guerrieri and Lorenzoni 2017; Bayer et al. 2019; Kaplan et al. 2018; Ahn et al. 2018; Auclert et al. ming). These insights have sparked a burgeoning body of literature, establishing HANK models as the "very model of modern monetary policy".[1]

This PhD thesis contributes to the ongoing advancements in the HANK literature in several ways. First, it employs empirically-realistic, medium-sized state-of-the-art HANK models to analyze contemporary policy issues such as climate change. Second, it deepens our understanding of the interplay between (systematic) monetary and fiscal policy within this framework and explore their potential substitutability. And third, it connects HANK models with the fast growing field of behavioral macroeconomics by integrating empirically-supported behavioral elements to enhance the empirical fit of HANK models and to address new puzzles and paradoxes emerging from the heterogeneous-agent revolution in monetary economics.

[1] www.imf.org/en/Publications/fandd/issues/2023/03/modern-monetary-policy-kaplan-moll-violante. The literature reviews of the individual chapters summarize large parts of this literature.

In the **first chapter**, I develop a Climate-HANK model by extending a medium-sized state-of-the-art HANK model by a carbon dioxide market. This allows me to analyze the business cycle and distributional impact of carbon pricing policies. By now there is a large consensus that carbon pricing should be part of the policy mix to fight climate change. Yet, carbon pricing introduces price fluctuations at business cycle frequencies which might have business cycle implications. Using my Climate-HANK model, I show that an adverse carbon shock increases inflation and reduces economic activity such as output, inflation and investment. I show that the economic costs of higher carbon prices are borne relatively evenly across the income and wealth distribution. Fiscal policy can mitigate the macroeconomic and welfare effects of higher carbon prices in a revenue-neutral way by using the extra revenue from higher carbon prices. If these revenues are paid back to households in the form of lump-sum transfers, higher carbon prices can even have a welfare-enhancing effect that goes beyond the welfare gain from lower emissions—a double dividend result.

The **second chapter**, which is joint work with Hannah Seidl, analyzes unconventional fiscal policy tools in HANK models and how they can be used to substitute for conventional monetary policy. In particular, we derive a perfect substitutability result between fiscal policy and monetary policy. We show in closed form that three simple conditions for consumption taxes, labor taxes, and the government debt level are sufficient to induce the same consumption and labor supply of each household and, thus, the same allocation as interest rate policies. The intuition is that consumption taxes and labor taxes replicate the effects that monetary policy has on the wedges in the first-order conditions of households. Combining these with debt policies can then in addition generate equivalence in the budget constraint of each household. Given our result, fiscal policy can replicate any allocation that hypothetically unconstrained monetary policy would generate when monetary policy is constrained by a binding lower bound, an exchange rate peg, or a monetary union.

How does a monetary union alter the impact of business cycle shocks at the household level? This fundamental question is a prime example which cannot be answered in models with representative agents. Thus, in the **third chapter**, Christian Bayer, Alexander Kriwoluzky, and Gernot Müller, and I develop a two-country HANK model (HANK^2) to analyze this question. We show in closed form that, first, when aggregated across both countries and households, macroeconomic dynamics are independent of whether there is a monetary union or not. Second, the overall effect of the shock for households in specific brackets of the wealth (and income) distribution, aggregated across countries of residence, does not depend on whether monetary union is in place or not. In other words, the monetary union itself does not shift the impact of the shock *vertically* across wealth classes. Instead, third, a monetary union shifts the impact of the shock *horizontally* across borders within the brackets of the wealth distribution. We then extend our model to a medium-sized HANK^2 model and calibrate it to the euro area. This exercise reveals that a monetary union alters the impact of shocks most strongly in the tails of the wealth distribution but leaves the middle class almost unaffected.

The **fourth chapter**, which is joint work with Oliver Pfäuti, bridges the HANK literature with the behavioral macroeconomics literature by developing a behavioral heterogeneous-agent New Keynesian model (behavioral HANK). The model features a New Keynesian core and we allow for household heterogeneity and incomplete markets as well as bounded rationality in the form of cognitive discounting. Under cognitive discounting, households' expectations underreact to macroeconomic news, consistent with what we find in survey data. We show that the behavioral HANK model accounts for recent important empirical findings on the transmission mechanisms and effectiveness of monetary policy. In particular, monetary policy affects household consumption to a large extent through indirect effects which tend to amplify the effects of conventional monetary policy on consumption as the incomes of households that exhibit higher marginal propensities to consume are more exposed to aggregate income fluctuations induced by monetary policy. Announcements of future monetary policy changes, in contrast, have relatively weak effects on current economic activity, and the economy remains stable at the effective lower bound. In contrast to existing models, we account for these facts simultaneously without having to rely on a specific monetary or fiscal policy. When abstracting from either household heterogeneity or bounded rationality the model fails to do so. Given its consistency with the empirical findings about the monetary transmission, we use the our model to revisit the policy implications of inflationary supply shocks. We uncover a novel amplification channel of supply shocks: Cognitive discounting and the unequal exposure of households as well as their interaction amplify the inflationary effects of supply shocks such that inflation increases 2.5 times as strongly as in models abstracting from these features.

While the fourth chapter deals with how households' expectations react to macroeconomic news, **the fifth chapter** (joint with Oliver Pfäuti and Jonathan Zinman) focuses on households' perception of their own income risk and the macroeconomic implications thereof. In particular, we show how heterogeneity within behavioral biases explains heterogeneity in households' saving behavior which in turn has important implications for macroeconomic fluctuations and fiscal policies. Using survey data, we show that households with lower cognitive skills are more likely to be overly optimistic with respect to their future financial situation, they are more likely to be hand-to-mouth, and they tend to be overconfident with respect to their cognitive skills. Motivated by these facts, we then introduce heterogeneity in cognitive skills and overconfidence into a HANK model and show that our model can account for our empirical findings. Overconfidence proves to be the key innovation, driving households to spend instead of precautionary save and producing empirically realistic wealth distributions and hand-to-mouth shares and MPCs across the income distribution. We show that accounting for heterogeneity in overconfidence has important implications for fiscal policy: Providing liquidity is less effective in bringing households away from the borrowing constraint, low-income benefits lead to less crowding out of self insurance, and the optimal level of public debt is substantially lower than in standard models.

Chapter 1

Climate-HANK: Carbon Shocks, Fiscal Policy, and Double Dividends

Abstract

Carbon pricing introduces price fluctuations at business cycle frequency. I develop a heterogeneous-agent New Keynesian model with a carbon market to analyze the effects of these fluctuations on the business cycle and their distributional impact. An adverse carbon shock increases inflation and decreases economic activity with the costs borne relatively evenly across the income and wealth distribution. Fiscal policy can mitigate the macroeconomic and welfare effects of higher carbon prices in a revenue-neutral way using the revenue from higher carbon prices. Paying it back as lump-sum transfers, higher carbon prices have a welfare-enhancing effect that goes beyond welfare gains from lower emissions—a double dividend result.

Keywords: Carbon shocks; Fiscal policy; HANK; Distributional impact of climate change; Heterogeneity; Inequality; Households; Double dividends
JEL-Classification: D31, E62, E64, Q58

1.1 Introduction

Climate change is one of the biggest challenges of our times. As carbon dioxide emissions are one of the main drivers for climate change, policymakers in many countries introduce costs of emitting carbon dioxide with the aim to reduce its demand by firms and households. The European Union (EU), for example, has introduced the European Union Emission Trading System (EU ETS) as a cap and trade system to manage the amount of carbon dioxide emitted in the EU. If regulated by the EU ETS, agents need to buy sufficient carbon dioxide certificates to cover their carbon dioxide emissions putting a price on carbon dioxide issuance. Figure 1.1 shows the price per ton of carbon dioxide use over the last 4 years in the EU. It demonstrates that the price for carbon dioxide use is highly volatile and that it fluctuates at business cycle frequency. This raises the questions: How does carbon price volatility affect the business cycle? Who bears the cost of carbon pricing? And which fiscal policy should be implemented to mitigate the costs of carbon price fluctuations?

To answer these questions, I build a *Climate-HANK* model: I extend a heterogeneous-agent New-Keynesian (HANK) model as in Bayer et al. (2020b) by a carbon market. Carbon dioxide enters the model two-fold: first, it is used as an input in the production function of firms alongside capital and labor. Second, households directly consume carbon dioxide and I allow for heterogeneity within the intensity of carbon dioxide in households' consumption baskets which I match to German micro data. I start by simulating the business cycle effects of a carbon shock, that is a surprise drop in the amount of carbon allowances. This increases the price per unit of carbon dioxide use. I find that this affects the economy similar to a cost-push shock: output, investment, and consumption decrease while inflation increases and average welfare in the economy decreases. I also show that the distributional effects are rather small: both along the income and along the wealth distribution, the costs of the carbon shock are borne relatively uniformly. While poorer households' welfare are hit harder by the recession, the consumption of richer households has a higher carbon dioxide intensity.

The adverse carbon shock increases the revenue of the government from selling carbon certificates. In the baseline, I assume it to be absorbed by lower debt in the short-run reflecting the lack of earmarking of carbon tax revenues for certain fiscal instruments in most countries. I show that earmarking these revenues for certain fiscal instruments could mitigate the aggregate effects of the carbon shock and eases its welfare burden. However, there is a trade-off between macroeconomic stabilization and welfare improvement depending on the fiscal instrument for which the extra revenue is earmarked: while paying a capital subsidy is most effective in curbing the recession, average welfare of households is lower than in the no-earmark case. In contrast, repaying the extra revenue directly to households in form of a lump-sum transfer mitigates the recession only slightly but it even increases average welfare above the status quo of no carbon shocks. In that sense, repayment by transfers is the HANK-version of the seminal double dividend result (see e.g. Goulder

Figure 1.1: The price of emissions allowance in the EU

Notes: Cost per ton of carbon dioxide use in €. Source: EMBER.

(1995)) as in this case, higher carbon pricing is even welfare-increasing. However, this comes at the cost of introducing much larger distributional consequences than the carbon shock itself as in this case, higher carbon prices redistribute from the income- and wealth-rich to the poor.

More in detail, in this paper I extend the medium-sized HANK model in Bayer et al. (2020b) by explicitly modelling the market for carbon dioxide, accounting for the carbon dioxide emissions of households and firms. Based on the EU-ETS, I assume that the total supply of carbon dioxide is exogenously set by the fiscal authority. I allow for heterogeneity in the carbon dioxide share of households' consumption baskets, in line with German micro data. The other model features are by now fairly standard. Financial markets are incomplete and households face idiosyncratic risk which is why they self-insure through savings in a liquid and illiquid asset. As a result, there is a non-degenerate distribution of income and wealth. Wages and prices are sticky through the standard way in New Keynesian models. I calibrate the economy based on data for Germany—the largest country in the EU—including the degree of heterogeneity in the carbon dioxide intensity in the consumption baskets found for German households. My modelling choices allow me to capture the heterogeneity in carbon dioxide intensity conditional on income for which I find that the carbon intensity tends to be lower for lower income households. At the same time, my model can capture the large heterogeneity even when conditioning on income which I find to be large in the data.

I simulate a surprise and temporary reduction in the carbon dioxide allowances. This increases the price of carbon dioxide and, thus, the costs in firms' production and the price of the consumption basket of households. I find that in response to this cost-push shock, output, investment, and consumption fall, while inflation increases. The consumption Gini only moves slightly, which shows the small distributional impact of the carbon shock. This is also reflected in the welfare impact of a carbon shock along the wealth and income distribution. The average welfare impact for all ten deciles of the income as well as for the wealth distribution is negative. Yet, I find the welfare impact of the carbon shock to be mitigated by its positive impact on the government budget. While this leads to lower government debt in the short run, lower taxes bring government debt eventually back to its steady state level. These lower taxes in the long-run, themselves, have a positive effect on welfare thereby mitigating the overall loss in welfare from higher carbon prices.

I then use the model to compute counterfactual analyses in which fiscal policy directly spends its extra revenue. More specifically, I consider three different fiscal instruments through which the government directly repays its extra revenue from the carbon system: lump-sum transfers to households, lower income taxes, and higher capital subsidies. Repayment through transfers mitigates the fall in output through carbon taxes as it stimulates private consumption. However, the dampening effects on output are quantitatively small as the higher transfers further decrease investment. I find that average welfare now increases rather than decreases in response to an adverse carbon shock. This uncovers a novel, HANK-type, double-dividend result: using the extra revenue from higher carbon prices can be welfare-beneficial if the extra revenue is used to finance higher transfers. The reason is that higher transfers decrease consumption inequality and provide partial insurance against households' idiosyncratic income risk. Hence, in addition to the unmodelled welfare gain of lower carbon emissions, the restructuring of government expenditures induced by higher carbon prices adds an additional welfare gain. Yet, I show that the higher average welfare comes at the cost of higher distributional impact of carbon shocks: both along the income and the wealth distribution, it is mainly the bottom 50% who benefit from higher transfers while the richest deciles are slightly worse off compared to the baseline model.

I then show that repayment through a direct decrease in income taxes substantially mitigates the recession, especially in terms of output and consumption. In addition, it increases average welfare, albeit at a smaller magnitude than repayment through transfers. Thus, I show that the classical double-dividend result—using the revenue from the carbon system to reduce distortionary taxes—holds in my business cycle context. Yet, again, the double-dividend result comes at the cost of increasing the distributional impact of carbon shocks, but to a much smaller degree than repayment through transfers would.

Finally, I show that repayment through capital subsidies mitigates the recession the most and even increases output on impact. They are also the most effective in mitigating the fall in investment. Yet, with repayment through capital subsidies, average welfare falls even stronger than in the baseline. At the same time, it amplifies the distributional impact of carbon shocks as it is the wealthy households that directly benefit from the capital subsidy.

The paper is structured as follows. In the remainder of this section, I discuss the related literature. Section 1.2 provides a summary of the model. Most of the details are relegated to the appendix. Instead, the exposition focuses on the carbon market. Section 1.3 presents details of the calibration of the model and Section 1.4 the results. The final section offers some conclusions.

Related literature. My paper relates closely to two strands of the literature. First, there is the recent surge of HANK models which are used to revisit the transmission of traditional business cycle shocks and business cycle policies starting with the influential study of Kaplan et al. (2018), but also, for instance, Auclert (2019) and Bayer, Lütticke, Pham-Dao, and Tjaden (2019). This framework lends itself naturally to the analysis of fiscal policy more broadly (Auclert, Rognlie, and Straub 2018; Bayer, Born, Luetticke, and Müller 2023; Pfäuti, Seyrich, and Zinman 2024) and in particular to the analysis of tax policies (Le Grand et al. 2021; Bhandari et al. 2021; Seidl and Seyrich 2023).

Second, my paper relates to the strand of literature that analyzes the macroeconomic impact of climate change and, in particular, to its sub-strand that studies business cycle implications of carbon prices (Känzig 2023; Metcalf 2019; Metcalf and Stock 2023; Bernard and Kichian 2021; Konradt and Weder 2021; McKibbin et al. 2017; Goulder and Hafstead 2017). In contrast to these papers, I build a general-equilibrium business-cycle model with heterogeneous agents and rich portfolio choices which allows me to jointly analyze the macroeconomic impact of carbon pricing and its distributional impact as well as to run counterfactuals analyzing alternative repayment structures of the extra revenue from higher carbon prices. Langot et al. (2023) is closest to mine. While I focus on German data and resort to a model with rich portfolio choice, they focus on France and on a HANK model without portfolio choice and without investment. Also, they do not compute the welfare impact of carbon pricing shocks. In addition, while they find in French data that the consumption baskets of lower-income households is more energy-intensive, I find in German data, that the consumption basket of lower-income households is less carbon dioxide intensive which implies that the partial equilibrium effect of higher carbon prices is progressive rather than regressive.

1.2 A Climate-HANK model

I evaluate the effects of temporary fluctuations in carbon pricing in a medium-scaled HANK model. In particular, I extend the HANK model in Bayer et al. (2020b) by a carbon market. The carbon market resembles the structure of the energy market in Bayer et al. (2023). The model accounts for carbon dioxide emissions in production and in household consumption. The latter represents direct carbon dioxide emissions by households which may be due to heating, gasoline or transport services. The model features incomplete financial markets and assets with different liquidity (bonds and capital). The model captures household heterogeneity with respect to income, wealth, portfolios, and carbon dioxide intensity of consumption. The following is a brief summary of the model, with a particular focus on the carbon market. A full description of the model can be found in the Appendix.

1.2.1 Summary of the model

Markets are incomplete and households face idiosyncratic, that is, household-specific, risks but are able to self-insure. They can do so using a liquid asset that can be traded every period and an illiquid asset (physical capital) which can only be traded subject to a friction. As a result, households are heterogeneous in terms of income and wealth. Households with little wealth or households whose wealth consists mainly of illiquid assets (e.g. houses) have a high propensity to consume out of disposable income and transfers.

Prices and wages are sticky, as is common in the New Keynesian literature. The model consists of a firm sector, a household sector, and a government sector. The firm sector consists of (a) perfectly competitive intermediate goods firms that produce intermediate goods using capital and labor thereby emitting carbon dioxide; (b) final goods firms that operate under monopolistic competition and produce differentiated final goods from homogeneous intermediate goods; (c) capital goods producers that transform consumer goods into capital; (d) labor intermediaries that produce labor services by combining differentiated labor from (e) unions that differentiate the raw labor provided by households. Pricing by final goods producers goods and wage setting by unions is subject to frictions à la Calvo (1983).

There is a continuum of households that consume final goods and emit carbon dioxide. Households earn income from supplying (raw) labor and capital to the labor and capital markets and from owning their firm sector, absorbing any rents arising from the market power of unions and final goods producers and from diminishing returns to scale in capital goods production.

The government sector comprises a monetary policy and a fiscal authority. The fiscal authority levies taxes on labor income and distributed profits, issues government bonds and sells carbon certificates. The fiscal authority also operates a targeted transfer system. Monetary policy sets the nominal interest rate in the economy using a Taylor rule, that is, it adjusts the interest rate to inflation.

1.2.2 Modelling carbon dioxide emissions

A distinct and novel feature of my analysis is to model carbon dioxide emissions in a HANK framework. The way I model carbon dioxide emissions in the economy is closely related to how Bayer et al. (2023) model energy usage in the economy. In the following, I provide details in this regard, first discussing sources of carbon dioxide issuance and then turning to the market for carbon certificates.

Carbon dioxide is issued in two different ways: First, producing the consumption good issues carbon dioxide such that carbon dioxide—along with labor and capital—is an input to the production of intermediate goods. Carbon dioxide can be substituted away in the production but this leads to efficiency losses in production. This allows me to capture how carbon prices affect the supply side of the economy through its effect on industrial production. Specifically, I assume intermediate goods Y_t are produced with the (nested)

CES production function:

$$Y_t = \left((1-a_P)^{\frac{1}{\sigma_P}} \; {Y_t^P}^{\frac{\sigma_P-1}{\sigma_P}} + a_P^{\frac{1}{\sigma_P}} \; \left(E_t^Y\right)^{\frac{\sigma_P-1}{\sigma_P}} \right)^{\frac{\sigma_P}{\sigma_P-1}}, \text{ where } Y_t^P = (u_t K_t^s)^{\alpha} \; {N_t}^{1-\alpha}. \tag{1.1}$$

As this expression shows, the intermediate good is made of a physical input, Y_t^P, which, in turn, combines capital, K_t, with capacity utilization u_t, and labor, N_t, on the one hand, and carbon dioxide, E_t^Y, on the other hand. The coefficient α is the capital share, the coefficient σ_P captures the (short-run) substitutability of carbon dioxide in the production process, and a_P is the carbon dioxide issuance during production in normal times.

Second, households directly emit carbon dioxide as part of their consumption basket. This can be thought of burning gasoline while driving a car, the carbon dioxide emissions caused by transport services, or carbon dioxide issued by heating or other energy sources.[2] In addition, I allow households to differ in their carbon dioxide intensity in consumption. Total consumption c_{it} of household i at time t consists of the physical consumption good c_{it}^P and carbon dioxide, E_{it}^C, again combined in a CES aggregator:

$$c_{it} = \left(\left(1-a_{it}^C\right)^{\frac{1}{\sigma_C}} \; {c_{it}^P}^{\frac{\sigma_C-1}{\sigma_C}} + {a_{it}^C}^{\frac{1}{\sigma_C}} \; \left(E_{it}^C\right)^{\frac{\sigma_C-1}{\sigma_C}} \right)^{\frac{\sigma_C}{\sigma_C-1}}. \tag{1.2}$$

Here σ_C represents the elasticity of substitution in consumption, that is, it measures the extent to which carbon dioxide can be substituted for physical consumption goods as relative prices fluctuate.

Households differ in the carbon dioxide intensity of their consumption baskets, captured by a_{it}^C. I assume that the long-run share of carbon dioxide in consumption varies exogenously across households and over time. The transitions from low to high and from high to low carbon dioxide intensity are random but related to the income state of the household. Concretely, I assume for the probability $\rho(h, a^C)$ to switch from one carbon dioxide type to the other the following functional form:

$$\rho(h, a^C) = \bar{\rho} + (\mathbb{I}_{a^C=a_H^C} - \mathbb{I}_{a^C=a_L^C})A(h) + \mathbb{I}_{a^C=a_L^C} B, \tag{1.3}$$

where A is a linear function of the human capital quintile and B is a constant that captures that it is in general more likely to remain a low carbon dioxide type.

This allows me to capture two key dimensions of heterogeneity in households' carbon dioxide share in the data: First, there is a strong positive correlation between the carbon dioxide emissions and household income. Second, there is a large dispersion in the carbon dioxide emissions even conditional on income. I will further discuss the large heterogeneity in Section 1.3. However, while I allow for transitions in carbon dioxide-intensity types, I

[2] In the terminology of Kuhn and Schlattmann (2024), households emit carbon dioxide both *directly* as well as *indirectly* with the latter being the carbon dioxide emission coming from the physical consumption good that households consume.

model type transitions as infrequent, so that the carbon dioxide intensity of the household is very persistent—in line with the fact that households preferences probably change very slowly.

The differences in carbon dioxide intensity also imply heterogeneity in inflation rates across households when carbon prices move. Since carbon dioxide is a component of household consumption, an increase in carbon prices raises the household price index and, all else equal, leads to a reduction in real income and, potentially, consumption. This effect is more pronounced for households with high carbon dioxide intensity than for households with low carbon dioxide intensity.

1.2.3 Demand for carbon certificates and market clearing

I assume that whenever firms or households issue carbon dioxide, they need to be in possession of a carbon certificate. In the spirit of the EU ETS, I assume that policy sets the amount of carbon certificates, E_t. For simplicity, I assume that the EU ETS covers all economic activity. This implies that there is a price on issuing carbon dioxide, p_t^C. As there are no other costs involved when issuing carbon dioxide, the price of issuance one unit of carbon dioxide equals the carbon price $p_t^E = p_t^C$. I further assume that the government sells the certificates such that the government receives all revenue from the carbon certificate system. The demand for carbon dioxide from the production sector satisfies the following first-order condition:

$$p_t^C = mc_t(1-a_P)^{\frac{1}{\sigma_P}} \left(\frac{Y_t}{E_t^Y}\right)^{\frac{1}{\sigma_P}}, \tag{1.4}$$

and the demand for carbon dioxide from households satisfy the type-specific first-order conditions:

$$E_{it}^C = \frac{1-a_{it}^C}{a_{it}^C}(p_t^C)^{-\sigma^C} c_{it}^P. \tag{1.5}$$

Total carbon dioxide issuance in the economy then equals the amount of carbon certificates:

$$E_t = E_t^C + E_t^Y. \tag{1.6}$$

I model a carbon shock by an exogenous and temporary decrease in carbon certificates. More precisely, I assume E_t to follow the following AR(1)-process:

$$E_t = (1-\rho_E)\bar{E} + \rho_E E_{t-1} + \varepsilon_t^E, \tag{1.7}$$

1.2.4 Carbon shocks and the role of fiscal policy

An increase in carbon pricing increases the fiscal surplus. There is a wide debate how to repay this fiscal surplus. In my baseline, I assume that there is no direct repayment reflecting the fact that in most countries, revenue through carbon pricing is not earmarked for certain expenditures or tax cuts. To analyze the potential usefulness of such earmarking, I then also analyze three counterfactual scenarios in which the extra revenue is earmarked for i) repayment through lump-sum transfers, ii) repayment through lower income taxes, and iii) repayment through capital subsidies.

Baseline. In the no repayment case, I assume that the extra tax income generated by higher carbon taxes enters the budget of the government without being earmarked for certain expenditures or tax cuts. Hence, in the short-run, higher revenue from higher carbon prices decreases government debt. In the long-run, lower income taxes bring back government debt to its steady state level resembling the idea that the carbon dioxide emissions revenue will eventually reduce the tax revenue that the government needs to collect.

$$\frac{\tau_t}{\bar{\tau}} = \left(\frac{\tau_{t-1}}{\bar{\tau}}\right)^{\rho_\tau} \left(\frac{B_{t+1}}{\bar{B}}\right)^{(1-\rho_\tau)\gamma_B^\tau}. \tag{1.8}$$

Repayment through lump-sum transfers. In this case, I assume that the extra revenue generated by the increase in higher carbon prices is directly redistributed to households via uniform lump-sum transfers. This implies:

$$Tr_t = p_{E,t}E_t - \bar{p_E}\bar{E}. \tag{1.9}$$

Repayment through income taxes. In this case, I assume that the extra revenue generated by the increase in higher carbon prices is earmarked for lower income taxes. In particular, I assume that income taxes are reduced by τ_t^C, with:

$$\tau_t^C = \frac{(p_{E,t}E_t - \bar{p_E}\bar{E})}{TB}, \tag{1.10}$$

where TB is the income tax base.

Repayment through capital subsidies. In this case, I assume that the extra revenue generated by the increase in higher carbon prices is earmarked for capital subsidies. In particular, I assume that capital subsides s_t^C are directly paid to firms per unit of capital such that:

$$s_t^C = \frac{(p_{E,t}E_t - \bar{p_E}\bar{E})}{K}. \tag{1.11}$$

In all three counterfactuals, I assume that all effects on the government budget which

Table 1.1: Calibration of the carbon sector

	Description	Value
σ_P	Elasticity of substitution in production	0.200
σ_C	Elasticity of substitution in consumption	0.200
a_P	Share of carbon dioxide in production	0.017
a_{CH}	Proportion of carbon dioxide in consumption: Type "high"	0.018
a_{CN}	Proportion of carbon dioxide in consumption: Type "low"	0.010
$\bar{\rho}$	Persistence of high carbon dioxide state at median income	0.966
A	Slope of probability to stay in low carbon dioxide state	0.005
B	Shift in probability to remain in low carbon dioxide state	0.010

are not due to changes in carbon revenues are absorbed by debt in the short-run with debt being brought back to steady state in the long-run by taxes.

1.3 Calibration

As for the model exposition, I mainly focus here on my calibration strategy for the carbon market. For the rest of the calibration, I only briefly sketch the calibration strategy leaving the details for Appendix A.2.

I calibrate the model economy to Germany for two reasons: first, it is the largest economy in the EU and, second, there is detailed data on households' consumption which allows me to determine the degree of heterogeneity in the carbon dioxide intensity of households' consumption baskets. When matching the income and wealth distribution in Germany, I follow the strategy employed in Bayer et al. (2023): A targeted transfer system mimics the German minimum benefits system which pays transfers to households with income below a certain threshold. I then set key parameter values in order to match the debt ratio, the capital ratio, the wealth Gini, the share of the 10% richest in total wealth, the share of the 50% poorest in total wealth, and the share of indebted households in Germany. I set the remaining parameters to values that have been established in business cycle analyses based on New Keynesian models. Appendix A.2 provides details on the calibration.

Given the focus on this paper, I report the key parameters related to the carbon dioxide sector in some detail, see Table 1.1. Specifically, I choose the carbon dioxide share, a^P, for the firm sector to match the steady-state carbon dioxide expenditure shares of 1.7% of production costs.[3] I set the elasticity of substitution in production σ_P to 0.2. This captures the limited substitutability of carbon dioxide especially in the short-run reflecting the energy substitutability found in (Bachmann et al. 2022). For the household sector, I set the elasticity of substitution to $\sigma_C = 0.2$ which is an intermediate value between Bachmann et al. (2022)'s value for the energy elasticity of households and the value in

[3] In 2022, German production emitted 573 millions tons of carbon dioxide. Assuming a carbon price of 100€—which amounts to the highest value in the EU ETS so far—this equals 57 billions € which amounts to 1.7% of the production costs.

Table 1.2: Emission of carbon dioxide of households

Emission of CO2 Data (Model) in tons per quartal									
Income quintiles				Expenditure quartiles					
		Mean		p25		p50		p75	
		D	M	D	M	D	M	D	M
I:	0-20%	0.72	0.79	0.39	0.47	0.59	0.58	0.93	0.86
II:	20-40%	1.17	1.03	0.65	0.56	1.04	0.79	1.51	1.22
III:	40-60%	1.54	1.33	0.92	0.74	1.39	1.13	1.95	1.63
VI:	60-80%	1.97	1.77	1.25	1.04	1.80	1.70	2.48	2.14
V:	80-100%	2.39	2.75	1.50	1.93	2.15	2.70	2.97	3.24
Targets: relative moment by income quintile									
Mean(I)/Mean(V)		**0.30**	**0.29**	p25(III)/p75(III)		**0.47**	**0.46**		
p25(I)/p75(V)		**0.13**	**0.14**						

Source: German Einkommens- und Verbrauchsstichprobe (EVS) 2018, own calculations. Income quintiles refer to household net incomes. Expenditure quartiles refer to the within-income-quintile carbon dioxide emission. Columns *D* refer to the data, *M* to the model. Targeted moments in bold.

Känzig (2023). To compute the share of carbon dioxide in households consumption baskets, I take expenditures on the three consumption categories that have by the far the highest carbon intensity:[4] fuel, transportation services, and energy consumption.[5] In addition, I follow Bayer et al. (2023) and set $\bar{\rho} = 0.97$ to match the average probability to switch carbon dioxide types to roughly 10%. This leaves me with four additional parameters to characterize the carbon dioxide intensity of household consumption: the carbon dioxide share in the consumption basket of high and low carbon intensive households (a_H^C, a_L^C) and the parameters (B, A) which govern the process that determines the carbon dioxide type of households, given in Equation (1.3) above. I set these parameters so that the average expenditure share on carbon dioxide amounts to 1.4%,[6] and to capture the dispersion of carbon dioxide emissions within and across incomes as shown in Table 1.2. Concretely, I match the following three targets: (i) the average increase in carbon dioxide emissions across income quintiles, (ii) the interquartile range within the median income quintile, (iii)

[4]This calibration strategy resembles the calibration strategy in Kuhn and Schlattmann (2024).

[5]More precisely, I compute how much an individual household spend for these three consumption categories. I then compute the carbon dioxide per € ratio for these three sectors indicating how much carbon dioxide is emitted in these sectors per € and multiply this with households expenditures in these sectors. I then aggregate households carbon dioxide emission over these three sectors.

[6]The average German household emits 1.55 ton carbon dioxide quarterly. Assuming again a current carbon price of 100€ per ton, this amounts to 155€ spent on carbon dioxide per household quarterly which represents 1.4% of average private consumption.

the bottom quartile of carbon dioxide emissions in the bottom income quintile relative to the top quartile of carbon dioxide emissions in the top income quintile.

In this way, I capture a high gradient of carbon dioxide emissions with income, that is, some non-homotheticity in carbon dioxide emissions on average, without resorting to non-homothetic preferences themselves. At the same time, I capture the large dispersion in carbon dioxide emissions even conditional on income. In fact, Table 1.2 also shows that the non-targeted carbon dioxide emissions (relative to the average) of the different groups in the carbon dioxide issuance and income distribution are relatively well matched, despite the very coarse parameterization. It shows the umbers implied by the model alongside the empirical distribution from German micro data (the German equivalent of the Consumption Expenditure Survey, CEX).

1.4 Results

In what follows, I study the effects of carbon shocks through the lens of the calibrated model. I compute a linearized state-space solution using the toolkit provided by Bayer et al. (2020b).

1.4.1 Baseline

Figure 1.2 shows the macroeconomic effects of an adverse carbon shock: the amount of carbon certificates decreases on impact by 2.5% and gradually converges back to steady state. This increases the carbon price by 25%. Consequently, output, consumption, and inflation decline by 0.22%, 0.29%, and 0.42% at their respective peaks. Year-on-year headline inflation increases by at most 0.4 percentage points. The average consumption of households with a high carbon dioxide intensity in their consumption basket falls by around 50% more than that of households with a low carbon dioxide intensity. Consumption inequality, measured by the consumption Gini, falls mildly. The higher carbon price increases the revenue from the carbon certificate trade which is absorbed by lower public debt in the short run. Taxes gradually decreases to bring back debt in the long-run to its steady state level.

Who bears the cost of higher carbon prices? To quantify the impact of carbon shocks at the household level, I measure the welfare impact using the consumption equivalent variation, which is the permanent consumption change that would make an individual household equally well off as the shock under consideration.[7] Figure 1.3 shows the impact of the adverse carbon shock on welfare along the income distribution (left panel) and along the wealth distribution (right panel). It depicts the average welfare of the respective income and wealth decile. The results document that higher carbon prices decrease average welfare

[7] I take an ex-post perspective, evaluating welfare based on the specific shock at hand (that is, one-sided welfare), rather than providing an ex-ante welfare analysis based on a second-order approximation of the utility function.

Figure 1.2: Response to an exogenous decrease in CO2 certificates, baseline

(1.2.1) Carbon certificates, E

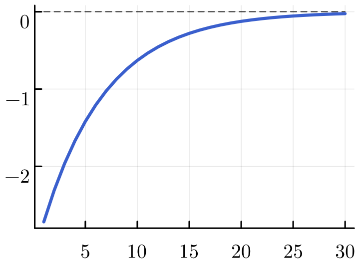

(1.2.2) Carbon price, p^E

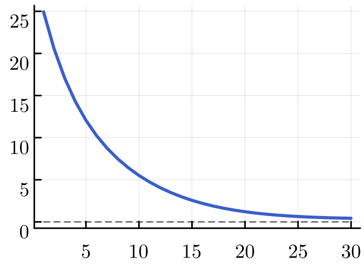

(1.2.3) Taxes, τ^L

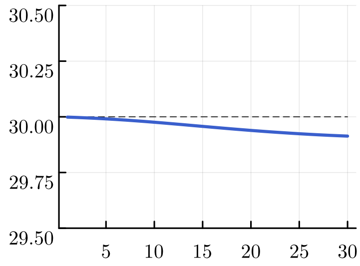

(1.2.4) Output, Y

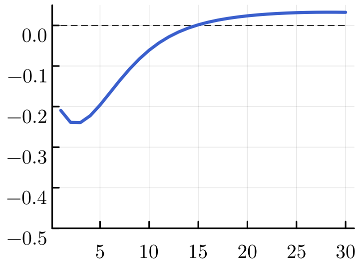

(1.2.5) Consumption, C

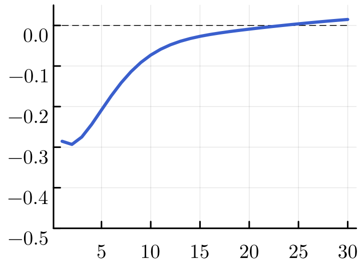

(1.2.6) Investment, I

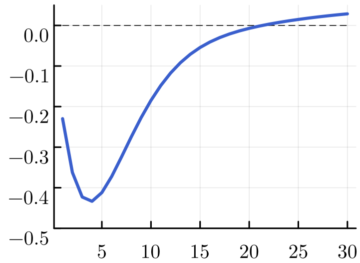

(1.2.7) Head inflation, π

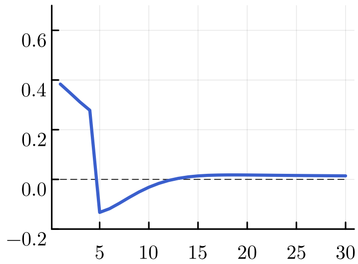

(1.2.8) Consumption Type Low, L

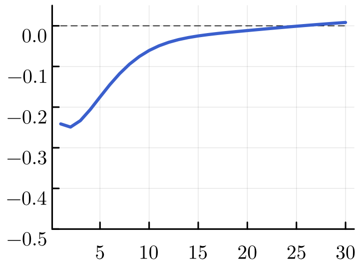

(1.2.9) Consumption Type High, H

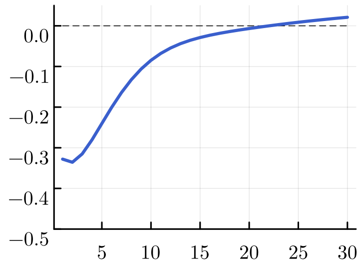

(1.2.10) Government debt, $\frac{B}{Y}$

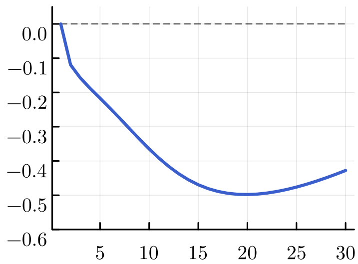

(1.2.11) Nominal Interest Rate, R^b

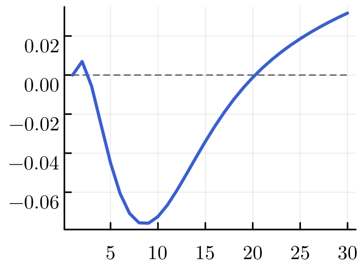

(1.2.12) Consumption Gini

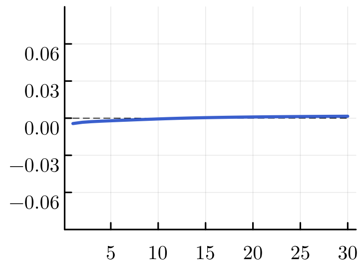

Impulse responses after an adverse carbon shock. Y-axis:Percentage deviation from steady state, percentage points in case of inflation and interest rate, and tax level in percentage in case of taxes. X-axis: Quarters.

for all deciles. There is some heterogeneity in the impact of higher carbon prices. Whether the general equilibrium effects of higher carbon prices affect the poor, the middle class or rich households depends on whether income or welfare is taken as a basis: along the income distribution, higher carbon prices hit the middle-class hardest, whereas along the wealth distribution, it affects both asset-poor and asset-rich households hardest. Overall, however, the costs of higher carbon prices is quite uniform reflected by the fact that the welfare impact has the same sign for all deciles.

To better understand the welfare impact of carbon shocks along the income and wealth distributions, I decompose the overall effect into its partial equilibrium effects. For this purpose, I exploit the fact that how the welfare of a given household is affected by a carbon shock depends on the arguments that enter its decision problem. I group the arguments into i) the carbon price itself, ii) changes in taxes, iii) changes in wage and profit income, and iv) changes in interest rates including the return on liquid bonds and the return on capital. These four partial equilibrium effects are reported in the lower four panels of Figure 1.3.

Figure 1.3: The welfare and distributional effects of carbon shocks

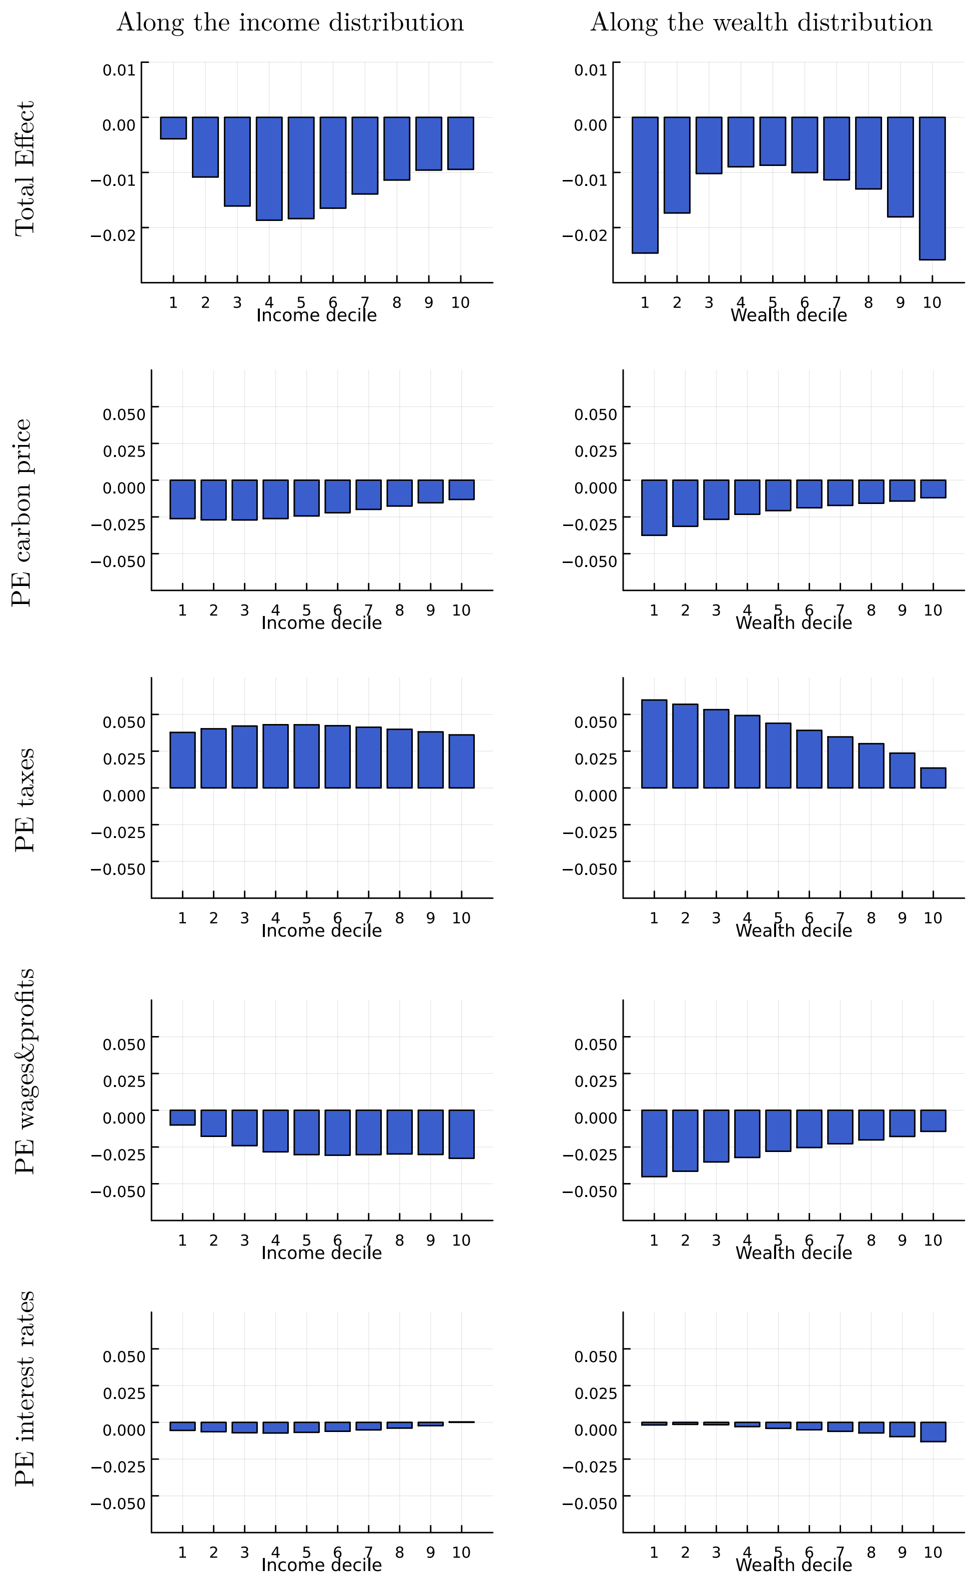

Notes: Welfare impact of a carbon shock along the income (left panels) and the wealth distribution. The upper panel shows the total impact whereas the lower panels decompose the total effects in its partial equilibrium components. Y-axis: Consumption equivalent compensating variations. X-axis: Income/wealth deciles.

There are three important observations from this exercise. First, focusing on the income distribution, the middle-income households experience the strongest impact from three of the four partial equilibrium effects.[8] This explains why middle-income households are most affected by the carbon shock. Second, focusing on the wealth distribution, households with the lowest wealth are the most affected, as they have the smallest buffer stock and are thus hit hardest by rising prices and falling incomes. The exception is the partial equilibrium effect of interest rates, which impacts the wealthiest households the most, though this partial equilibrium effect is quantitatively small. Third, all partial equilibrium effects reduce welfare in each income or wealth decile, with the exception of the partial equilibrium effect of taxes. Taxes have a positive and a rather large impact on the welfare of households. This demonstrates the "double dividend" of higher carbon prices as they allow the government to decrease distortionary taxes. Similar to the classical double-dividend result (see, e.g., Goulder (1995)), the eventual reduction in distortionary taxes is insufficient to prevent higher carbon prices from being detrimental to welfare (excluding the welfare gains from reduced carbon dioxide emissions from which I abstract in this paper). However, it does account for why the overall impact on welfare is relatively moderate: the average consumption equivalent is only -0.015%.[9]

1.4.2 Counterfactual scenarios: immediate repayments of carbon revenues

So far, I have assumed that fiscal policy does not respond to changes in carbon revenues, reflecting the status quo in many countries of not earmarking carbon revenues for specific fiscal instruments. As shown in the previous section, an increase in carbon prices increase the revenue from selling carbon certificates and, thus, reduces government debt quite substantially. What if fiscal policy immediately repays this extra revenue? To answer this question, I simulate three counterfactual scenarios: I assume that all extra carbon revenues are immediately i) paid back by lump-sum transfers to households, ii) used to decrease income taxes, iii) used to pay a capital subsidy.

Figure 1.4 shows the macroeconomic effects of the adverse carbon shock in the baseline (left panels) and in the three counterfactuals. To economize on space, I only show impulse responses for output, consumption, investment, inflation, and the consumption Gini.[10] In all three counterfactual cases, the recession caused by an adverse carbon shock is smaller than in the baseline, but the extents of mitigation differ across the counterfactual cases: While in the transfer counterfactual, output is only slightly different from the baseline, it falls not

[8] This is not the case for wage & profit income which affects the high-income households the most.

[9] This small loss in welfare heavily depends on the fact that the higher carbon revenue eventually leads to lower distortionary taxes. If instead, the higher revenue would be used for inefficient "rent-seeking", the welfare impact would be much higher. I approximate this scenario assuming that government debt will be brought back to steady state by higher wasteful government spending in the long-run. In this case, the average consumption equivalent is -0.060% and, thus, the welfare loss would be four times as high.

[10] The full set of impulse responses for the counterfactuals can be found in Appendix A.3.

even half as strong in the tax counterfactual. And in the capital subsidy counterfactual, output actually increases rather than decreases on impact. Non-surprisingly, capital subsides are also most effective in mitigating the fall in investment. In contrast, investment falls even more in the transfer counterfactual than it does in the baseline.

The carbon shock is slightly more inflationary in the transfer counterfactual than in the baseline and even much more inflationary in the capital subsidy counterfactual. In contrast, the carbon shock affects inflation the same in the tax counterfactual as it does in the baseline.[11] In all three counterfactuals, the consumption Gini responds much more strongly than in the baseline, in which it remains almost constant. The consumption Gini also differs strongly across the three different counterfactual scenarios: While consumption inequality falls in the tax case and even more strongly in the transfer case, it strongly increases in the capital subsidy case. The latter reflects the fact that mostly wealthy households benefit from capital subsidies.

[11] In the HANK model, income taxes have both a supply side effect (through lower wages and hence lower marginal costs) as well as a demand side effect (through higher disposable income). These two opposite effects on inflation cancel out, such that they do not affect the inflation response of carbon shocks.

Macro effects

Figure 1.4: The macro effects of carbon shocks, alternative repayment methods

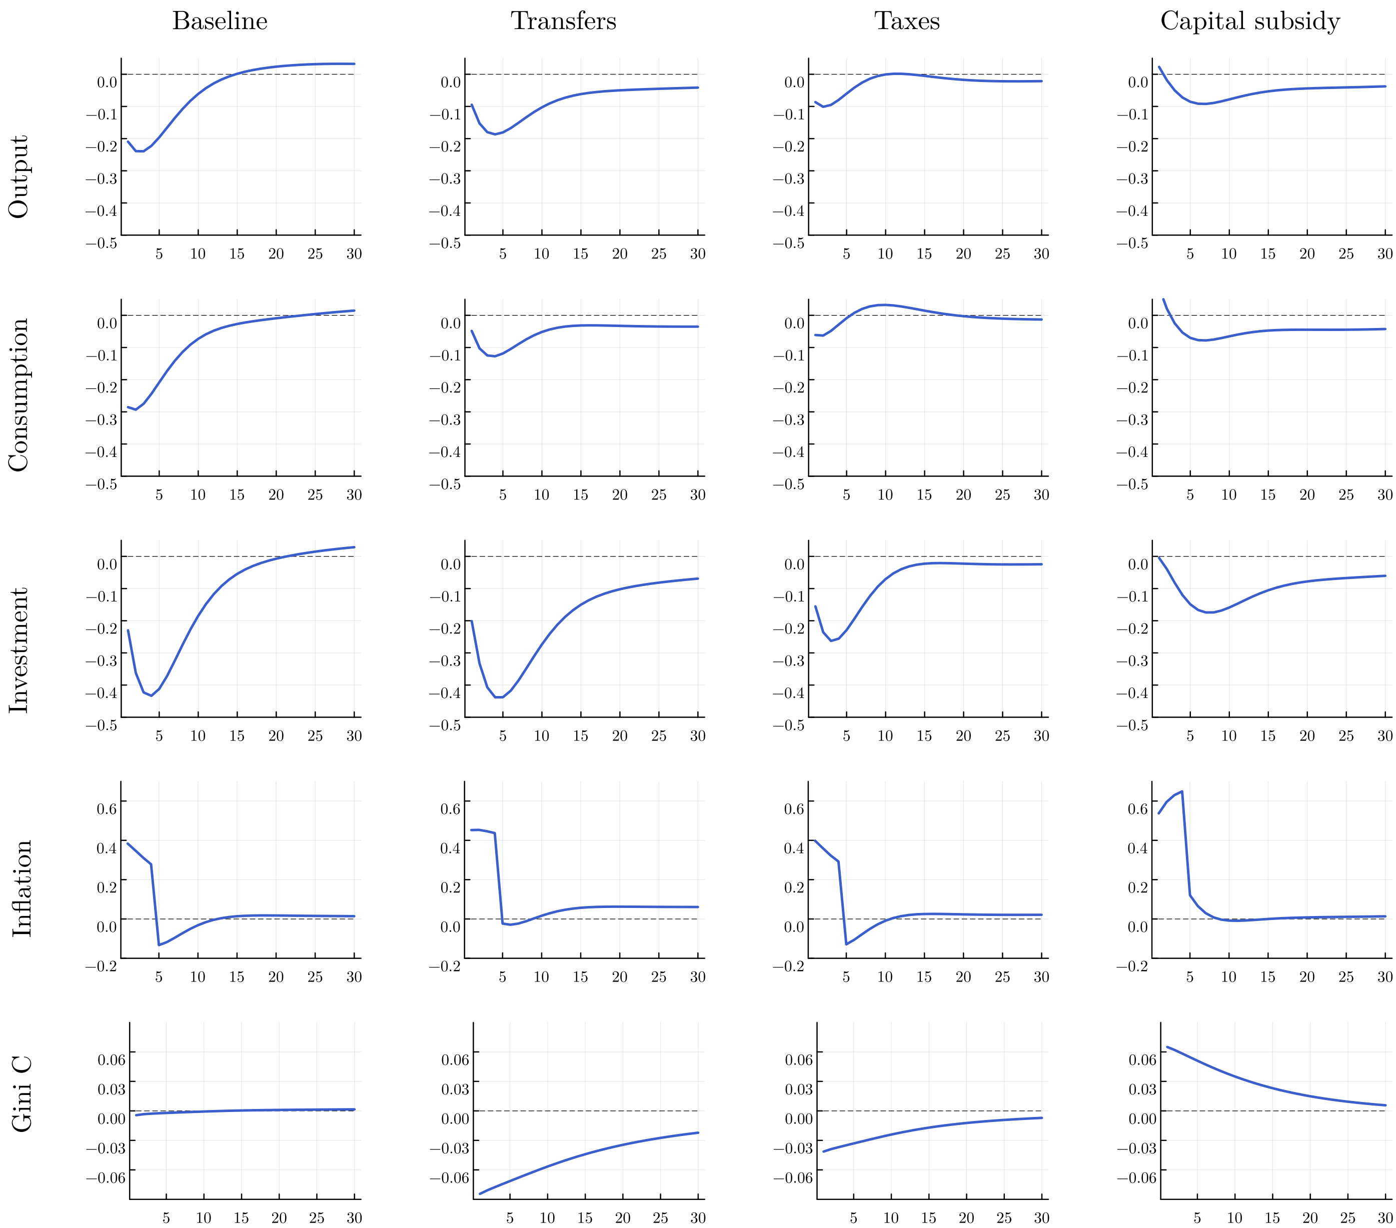

Notes: Impulse responses after an adverse carbon shock. Y-axis:Percentage deviation from steady state, percentage points in case of inflation. X-axis: Quarters.

Distributional effects

Figure 1.5 compares the welfare impact of the three counterfactual cases (lower panels) with the baseline case (upper panel). It does so again along the income distribution (left panel) and along the wealth distribution (right panel). All three counterfactuals changes the welfare impact of carbon shocks quite drastically. First, in all counterfactuals, it is no longer the case that the welfare impact of each income and wealth decile has the same sign. In other words, there are now winners and losers from the higher carbon prices amplifying the distributional impact of carbon pricing. In the transfer and in the tax case, the poorer households benefit from higher carbon prices whereas richer households suffer. This distributional differences are particularly strong in the transfer case as the induced higher lump-sum transfers redistribute more strongly towards relatively poorer households. In the capital subsidy case, most of the households suffer from carbon prices, with the exception of the wealthiest households which benefit from it. This reflects the fact that capital subsidies redistributes towards asset-rich households.

Second, for the majorities of deciles, the welfare impact is stronger in the counterfactual cases than in the baseline case. This implies that the mitigated macroeconomic effects in the counterfactual cases depicted in Figure 1.4 actually hide stronger adjustments on the micro level.

The third observation regards the *average* sign and size of the welfare impact. In the tax and in the transfer case, average welfare actually increases rather than decreases after an adverse carbon shock. This adds a second dividend to the unmodelled gain in welfare due to lower carbon emissions. This effect is particularly pronounced for the transfer counterfactual in which case the average consumption equivalent is 0.040% (whereas it is 0.008% in the tax case). Paying the higher carbon revenue to households via lump-sum transfers decreases consumption inequality and partly insures households' idiosyncratic income risk. This is sufficiently welfare-enhancing to dominate the detrimental welfare costs of higher carbon prices which can be seen as a HANK-version of the double-dividend result.

In contrast, average welfare falls even more strongly in the capital subsidy case with an average consumption equivalent of -0.046% instead of -0.015% in the baseline. While capital subsidies are effective in preventing output from falling on impact, they prevent the eventual fall in distortionary taxes in the long-run which occurs in the baseline case. This highlights a trade-off between macroeconomic stabilization and welfare impact when it comes to the choice of fiscal instrument for repaying higher carbon revenues: On the one hand, the instrument which is most effective in stabilizing output on impact (capital subsidies), increases the welfare loss of higher carbon prices even further. And on the other hand, the instrument that is best for average welfare (lump-sum transfers) is not very effective in mitigating the recession.

Figure 1.5: The distributional effects of carbon shocks, alternative repayment methods

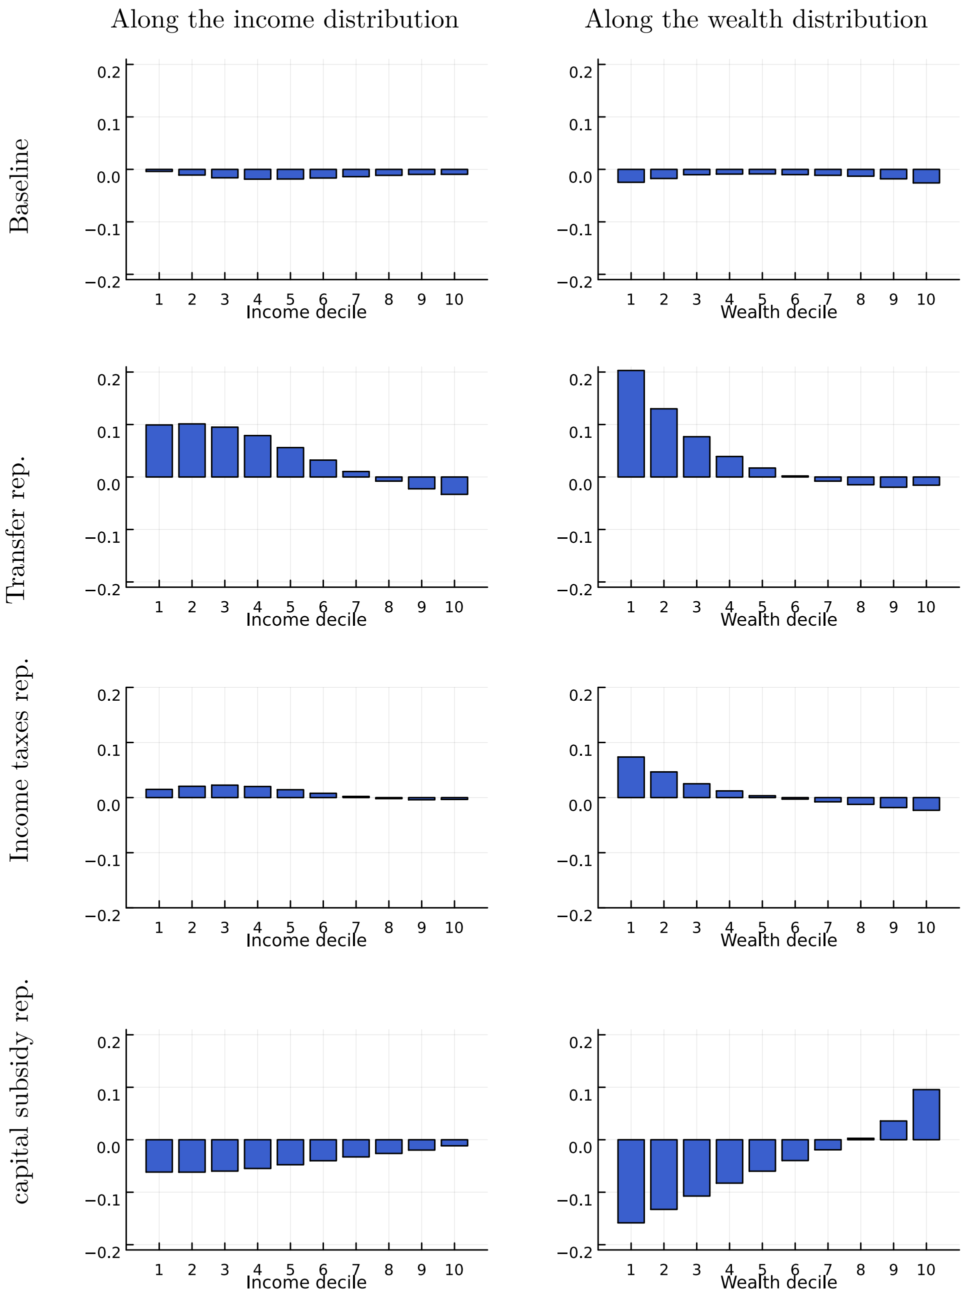

Notes: Welfare impact of an adverse carbon shocks along the income distribution (left panels) and along the wealth distribution (right panels). Y-axis:Consumption equivalent compensating variations. X-axis: Income/wealth deciles.

1.5 Conclusion

There is a consensus that combating climate change requires assigning a price to carbon dioxide emissions. However, implementing such pricing systems can be a source of business cycle fluctuations. This paper introduces a Climate-HANK model to study the business cycle impact and distributional effects of carbon shocks. The findings indicate that adverse carbon shocks influence the economy similarly to cost-push shocks, leading to reduced economic activity and increased inflation. Yet, I find the distributional impact of these shocks to be quantitatively small.

Although higher carbon prices can induce recessions, they can still increase overall welfare. The reason is that higher carbon prices increase the revenue of the government. I show that if the government pays this extra revenue directly back to households via lump-sum transfers—a repayment approach that has gained considerable support in recent years—higher carbon prices increase average welfare beyond the welfare gains from lower carbon emissions. This result represents a HANK adaptation of the classical double-dividend result in the environmental taxation literature. However, lump-sum transfers are not as effective in mitigating the macroeconomic impact of carbon shocks as other fiscal instruments. Future research could therefore address the question of which revenue-neutral fiscal mix is best suited to counteract carbon shocks, considering their macroeconomic, welfare, and distributional impacts.

Chapter 2

Unconventional Fiscal Policy in a Heterogeneous-Agent New Keynesian Model

with Hannah Seidl

Abstract

We show that in a New Keynesian model with household heterogeneity, fiscal policy can be a perfect substitute for monetary policy: three simple conditions for consumption taxes, labor taxes, and the government debt level are sufficient to induce the same consumption and labor supply of each household and, thus, the same allocation as interest rate policies. When monetary policy is constrained by a binding lower bound, a currency union, or an exchange rate peg, fiscal policy can therefore replicate any allocation that hypothetically unconstrained monetary policy would generate.

Keywords: Unconventional Fiscal Policy; Heterogeneous Agents; Incomplete Markets; Liquidity Trap; Sticky Prices
JEL-Classification: E12, E21, E24, E43, E52
This version was published as
Seidl, H. and F. Seyrich (2023). Unconventional Fiscal Policy in a Heterogeneous-Agent New Keynesian Model *Journal of Political Economy Macroeconomics* 1.4, 633-664. doi: `https://doi.org/10.1086/727506`

2.1 Introduction

Monetary policy oftentimes cannot freely adjust the nominal interest rate—be it due to a binding lower bound, a currency union, or an exchange rate peg. In these environments of constrained monetary policy, policymakers need to resort to alternative stabilization tools. Recent real-world episodes suggest that unconventional fiscal policy tools such as changes in the consumption tax rates are promising alternatives to conventional interest rate policies as they stimulate consumption through the intertemporal substitution channel (Bachmann et al. (2021), Baker et al. (2019), D'Acunto et al. (2022)).

In this paper, we show that a mix of unconventional fiscal policy tools can be a perfect substitute to monetary policy in a heterogeneous agent New Keynesian (HANK) model. In particular, we show that three simple conditions for consumption taxes, labor taxes, and the government debt level are sufficient to generate the same consumption and labor supply of *each* household and, thus, the same allocation as monetary policy. This perfect substitutability result holds when monetary policy is constrained meaning that our unconventional fiscal policy measure—which we label *HANK unconventional fiscal policy (HANK-UFP)*—circumvents the constraints of monetary policy. In particular, we show at the Effective Lower Bound (ELB) that HANK-UFP can generate any allocation that hypothetically unconstrained monetary policy could achieve by inducing the same cross-sectional consumption and labor supply.

The intuition for our perfect substitutability result is that HANK-UFP affects the optimization problem of each household in the same way as a change in the interest rate thereby replicating its whole transmission mechanism: HANK-UFP and interest rate changes induce the same inter- and intratemporal incentives for consumption and labor supply as well as the same effects on each household's budget constraint. In that sense, our analysis builds on the perfect substitutability result between fiscal and monetary policy in a representative agent New Keynesian (RANK) model (see Correia et al. (2008), Correia et al. (2013)). In RANK, consumption taxes and labor taxes alone—*unconventional fiscal policy (UFP)*—are sufficient to induce the same optimization problem of the representative household as monetary policy since they induce the same inter- and intratemporal incentives for consumption and labor supply. However, this result relies on the fact that, by construction, policies do not redistribute across households as all of the income accrues to the same household. In contrast, in HANK models, households are heterogeneous both in their income and in their income compositions. One of our contributions is to show that, as a consequence, tax policies alone are no longer sufficient to induce the same optimization problem of each household in HANK. The reason is that tax policies alone have different effects on the various income components of households and, hence, affect households' budget constraints differently than interest rate policies. In addition, these different cross-sectional effects induce different aggregate effects since households are heterogeneous in their marginal propensities to consume (MPCs). We show that, as a consequence, tax policies alone are not able to stabilize the economy while the ELB is binding. Furthermore, tax policies alone push the economy to a new steady state after the ELB stops binding characterized by lower real interest rates and higher inefficiencies out of incomplete markets compared to the original steady state.

For our analysis, we extend the textbook New Keynesian model by a standard heterogeneous agent, incomplete markets set-up. We assume that households face uninsurable idiosyncratic income risk and borrowing constraints. Households self-insure against idiosyncratic shocks to their labor productivity by buying risk-free bonds. Monetary policy sets the interest rate and fiscal policy sets proportional taxes on consumption and on labor, issues government debt, and pays lump-sum transfers to households. We assume that households have perfect foresight. In this environment, we analytically characterize our fiscal policy scheme, HANK-UFP, which induces the same optimization problem of each household as a change in the interest rate. This is sufficient to generate the same effects through general equilibrium since monetary policy and fiscal policy do not affect firms' equilibrium conditions directly but only through the household block.

To fix ideas, consider the effects of expansionary monetary policy on the households' optimization problems which are replicated by HANK-UFP as follows. In line with Correia et al. (2013), tax policies generate the same inter- and intramteporal incentives as monetary policy: pre-announced paths for higher future consumption taxes trigger the same incentives to intertemporally substitute consumption as a decrease in the real interest rate. Yet, higher consumption taxes also incentivize households to reduce their labor supply for any given real wage. Lower labor taxes offset this effect of consumption taxes on the labor supply.

We prove that when these tax policies are combined with debt policies in the form of an increase in the government debt level, each household's income and therefore her budget constraint is identically affected by HANK-UFP and monetary policy. We show that for this to be the case, it is sufficient that monetary policy and HANK-UFP induce the same redistribution through the policy block.[12] Expansionary monetary policy redistributes from asset holders to the government: on the one hand, lower interest rates induce a negative wealth effect which affects households in proportion to their asset holdings. On the other hand, the government issues the assets and, hence, has lower interest rate payments which shifts resources to the government. These additional resources are then redistributed back to households through a fiscal response. HANK-UFP replicates this redistribution through the policy block as follows. Higher consumption taxes generate the same negative wealth effect on the assets of households as they decrease the consumption value of assets. This again hurts households in proportion to their asset holdings. As households accumulate

[12]Among others, Bhandari et al. (2021), Bilbiie (2021), and Acharya and Dogra (2020) highlight the effects of households' heterogeneous exposure to a policy change arising indirectly through changes in output. While households are also heterogeneously exposed to changes in output in our model, this does not affect our perfect substitutability result because these effects are the same with HANK-UFP and monetary policy since both identically affect output. Thus, for our analysis, it is sufficient to focus on the heterogeneous exposure of households to policy changes arising directly from changes in monetary and fiscal policy variables, that is, through the policy block.

these assets for self-insurance purposes, higher consumption taxes increase the precautionary savings demand of households in proportion to their asset holdings. The government accommodates this higher asset demand by increasing the government debt level such that the value of total assets in consumption value terms is the same as in the monetary policy case. This provides the government with the same additional resources as in the monetary policy case which triggers the same fiscal response and, thus, the same redistribution back to households.

Our perfect substitutability result between fiscal policy and monetary policy is especially relevant when conventional monetary policy is constrained. We therefore apply our perfect substitutability result at the ELB—a typical case of constrained monetary policy—and show that HANK-UFP circumvents the constraint. By increasing consumption taxes, decreasing labor taxes, and permanently increasing the government debt level, HANK-UFP replicates the allocation associated with hypothetically unconstrained monetary policy—the counterfactual in which monetary policy could freely set nominal interest rates without any lower bound constraints.

In this ELB environment, we quantify the role of debt policies—the novel instrument that is necessary for perfect substitutability in HANK. To this end, we study through the lens of our HANK model the UFP scheme of Correia et al. (2013) which only consists of consumption taxes and labor taxes replicating the inter- and intratemporal incentives of households. Since government debt is now not adjusted for the consumption value but the same as in the unconstrained monetary policy case, the fiscal response that monetary policy induces through reducing the interest rate payments cannot be replicated by this fiscal policy scheme. As a consequence, this tax policy scheme cannot fully stabilize the economy in the short-run. We further show that with this tax-only policy scheme the economy now converges to a new steady state in which the real interest rate is lower than in the original steady state. The reason is that fiscal policy cannot satisfy the higher precautionary savings demand of households in the long-run induced by the permanently higher consumption taxes. Consequently, households are worse insured against their idiosyncratic income risk which permanently increases the inefficiency from incomplete markets. We show that as a consequence, each household's welfare is lower compared to the HANK-UFP scenario. This shows that tax policies which are neutral in RANK in the long-run can induce long-run inefficiencies in HANK.

While our baseline model abstracts from sticky wages and capital, we show analytically that fiscal policy can also be a perfect substitute for monetary policy if we extend our model by these model features. We first include sticky wages implemented via unions into the model and show that the same three conditions for consumption taxes, labor taxes, and the government debt level are sufficient for perfect substitutability. The reason is that given HANK-UFP, the unions face the same wage-setting problem as in the monetary policy case since HANK-UFP replicates the consumption of each household. When we add capital to our model, monetary policy affects the return on capital which changes the incentives to use capital for firms and which affects households' budget constraints. We show that HANK-UFP can still be a perfect substitute for monetary policy if it is extended. A condition for capital subsidies paid to the firms allows HANK-UFP to replicate firms' incentives to use capital in the spirit of Correia et al. (2013). How HANK-UFP replicates the effects on households' budget constraints depends on the degree of substitutability of bonds and capital: if they are perfect substitutes, the effects on the budget constraints are replicated by an extra increase in the government debt level. If capital is illiquid and, hence, an imperfect substitute for bonds, the government needs to issue an additional asset which has the same pay-off structure as capital. In both cases, the intuition from our baseline result carries over in the sense that an increase in consumption taxes decreases the consumption value of capital holdings. Accordingly, the extra bonds or, in the illiquid capital case, the additional assets are held by households in proportion to their capital holdings. In both cases, these extra asset holdings, in turn, finance the capital subsidies. This way, HANK-UFP redistributes from capital-rich households to firms in the same way as lower interest rates redistribute from capital-rich households to firms and, thus, the effects on each household's budget constraint are the same.

Related literature. Feldstein (2002) and Hall (2011) propose to increase future consumption taxes when monetary policy is constrained by the ELB. Baker et al. (2019), Bachmann et al. (2021), and D'Acunto et al. (2022) empirically show that recent real-world episodes of consumption tax policies have stimulated consumption and, thus, aggregate output through the intertemporal substitution channel. We show how these consumption tax policies can be part of a larger fiscal mix which does not only replicate the intertemporal substitution channel but the whole transmission channel of monetary policy.

Correia et al. (2008) and Correia et al. (2013) show that a combination of consumption taxes and labor taxes is a perfect substitute for monetary policy in RANK by replicating its effects on the policy wedges in the household's first-order conditions. Bianchi-Vimercati et al. (2021) show that in RANK, the effectiveness of these tax policies is not affected by bounded rationality in the form of level-k thinking. We depart from the textbook RANK model in a different way than Bianchi-Vimercati et al. (2021) and show that Correia et al. (2013)'s seminal result relies on the fact that monetary and fiscal policy do not redistribute among households in RANK. We show that when households are heterogeneous in their income composition, UFP as prescribed by Correia et al. (2013) is no longer a perfect substitute for monetary policy. We further show that in HANK, tax policies alone cannot fully stabilize the economy in the short-run and, in addition, push the economy to a new steady state characterized by a lower real interest rate.

Our analysis is also related to a large literature on the transmission mechanism of monetary policy in HANK (see among many others Werning (2015), McKay et al. (2016), Kaplan et al. (2018), Bilbiie (2021), Auclert (2019), Hagedorn et al. (2019), Acharya and

Dogra (2020), Auclert et al. (2020), Luetticke (2021)).[13] Recently, the HANK literature has also studied fiscal policy. Auclert et al. (ming), Ferriere and Navarro (2022), and Hagedorn et al. (2019) analyze fiscal multipliers in HANK models. Unlike our paper, these papers do not study whether fiscal policy can replicate the allocations of unconstrained monetary policy. Oh and Reis (2012) and Bayer et al. (2020) show that transfer policies can have large effects in HANK models which is also reflected in our numerical analysis. As in our analysis, Bayer et al. (2023) show that the government debt level affects the economy in the short- and in the long-run. In particular, they show how the government debt level affects the liquidity spread which, in turn, affects the real economy. In contrast, we provide a specific rule for the government debt level as part of a set of sufficient conditions such that fiscal policy is a perfect substitute for monetary policy. Bhandari et al. (2021) analyze optimal fiscal and monetary policy in a HANK model with aggregate risk. They analyze how idiosyncratic insurance possibilities shape optimal monetary and fiscal policy in the face of aggregate shocks while we focus on the substitutability of monetary policy by fiscal policy. Unlike our analysis, Bhandari et al. (2021) include aggregate risk, but they abstract from borrowing constraints, a binding ELB, and consumption taxes. Le Grand et al. (2021) also study optimal fiscal and monetary policy in a HANK model. Their set of fiscal instruments can replicate the allocations of monetary policy away from the ELB in their HANK model. Unlike our HANK-UFP scheme, they use capital taxes instead of consumption taxes. This is an important distinction from our analysis for two reasons: first, capital taxes face the same limitations as interest rate policies since the nominal after-tax return on savings cannot become negative. Hence, capital taxes cannot be used to circumvent the ELB constraint. Second, capital taxes have different effects on the budget constraints of households compared to consumption taxes.

Wolf (2021) is closest to our paper. He shows that transfer policies can achieve the same aggregate outcomes as interest rate cuts in a linearized HANK model with sticky wages. Unlike our perfect substitutability result, Wolf (2021) derives his aggregate equivalence result by showing that lump-sum transfers can trigger the same *aggregate* consumption response in partial equilibrium as monetary policy. In a linearized environment with sticky wages in which labor unions consider the marginal utility of aggregate consumption when setting wages, this then generates the same responses of macroeconomic aggregates through general equilibrium. Unlike Wolf (2021), we show that with tax and debt policies, fiscal policy can directly manipulate the optimization problem of each household in the same way as monetary policy. Thus, our result differs from the result in Wolf (2021) in two aspects: first, HANK-UFP does not only achieve equivalence in aggregates but also in the cross-section of households and, thus, the distribution of households evolves in the same way with monetary and fiscal policy. Second, this implies that our result also holds

[13] There is also a growing literature analyzing the transmission mechanism of monetary policy in models with firm heterogeneity, see among others Reiter et al. (2013), Koby and Wolf (2020), and Ottonello and Winberry (2020).

if households' labor supply is affected by their individual consumption in the short-run as the cross-sectional consumption and labor supply is the same as with monetary policy. To the best of our knowledge, we are the first who show how fiscal policy can circumvent constraints of monetary policy, including the ELB constraint, in a HANK model.

Outline. Section 2.2 presents our HANK model. Section 2.3 shows analytically that HANK-UFP is a perfect substitute for monetary policy. Section 2.4 provides a numerical analysis to show how HANK-UFP circumvents the ELB constraint and highlights the role of debt policies in HANK-UFP. Section 2.5 concludes.

2.2 Model

This section outlines our HANK model which is a sticky-price New Keynesian model extended by a standard heterogeneous households, incomplete markets set-up.

2.2.1 Households

The economy is populated by a continuum of households who are identical in their preferences given by:

$$E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{c_{h,t}^{1-\gamma}}{1-\gamma} - \frac{l_{h,t}^{1+\psi}}{1+\psi} \right], \tag{2.1}$$

where β denotes the household's discount factor, $c_{h,t}$ denotes consumption of household h in period t, and $l_{h,t}$ denotes her labor supply. The parameters γ and ψ govern the degree of risk aversion and the inverse Frisch elasticity, respectively.

The budget constraint of household h and her borrowing constraint are given by:

$$\begin{aligned}(1+\tau_t^C)c_{h,t} + b_{h,t+1} = (1+r_t)b_{h,t} + (1-\tau_t^L)w_t z_{h,t} l_{h,t} + D_t + Tr_t \\ b_{h,t+1} \geq 0,\end{aligned}$$

where $c_{h,t}$ denotes consumption of household h, $b_{h,t}$ are 1-period risk-free bonds which are issued by the government. r_t is the real interest rate paid on these bonds between period $t-1$ and period t. In addition, there is a proportional tax rate on consumption, τ_t^C, and a proportional tax rate on her individual labor income, τ_t^L. The labor income consists of the wage rate, w_t[14], the individual productivity level, z_t, and the individual labor supply. Since $z_{h,t}$ evolves according to an exogenous finite-state Markov chain, households face idiosyncratic income risk. As in McKay et al. (2016), we assume that all households receive

[14] Unless stated otherwise, all variables are denoted in real terms.

an equal share of firms' dividends, D_t[15], and a lump-sum transfer, Tr_t, from the government. For our analytical analysis in Section 2.3, it is useful to represent the budget constraint in consumption value terms:

$$c_{h,t} + \frac{b_{h,t+1}}{1+\tau_t^C} = (1+r_t)\frac{b_{h,t}}{1+\tau_t^C} + \frac{1-\tau_t^L}{1+\tau_t^C} w_t z_{h,t} l_{h,t} + \frac{D_t + Tr_t}{1+\tau_t^C} \tag{2.2}$$

$$b_{h,t+1} \geq 0.$$

We assume the standard Bewley-Huggett-Aiyagari incomplete markets setup such that there are no state-contingent securities. As households cannot buy perfect insurance, they accumulate government bonds to self-insure their idiosyncratic risk. As a consequence, households differ in their individual states which consists of household's asset position, b, and her specific productivity level, z. The decision problem of a household h is given by:

$$V_t(b_{h,t}, z_{h,t}) = \max_{c_{h,t}, l_{h,t}, b_{h,t+1}} \left\{ \frac{c_{h,t}^{1-\gamma}}{1-\gamma} - \frac{l_{h,t}^{1+\psi}}{1+\psi} + \beta \sum_{z_{h,t+1}} Pr(z_{h,t+1}|z_{h,t}) V_{t+1}(b_{h,t+1}, z_{h,t+1}) \right\},$$

subject to equation (2.2). The Euler equation is given by:

$$c_{h,t}^{-\gamma} \geq \beta E_t \left\{ \left((1+r_{t+1}) \frac{1+\tau_t^C}{1+\tau_{t+1}^C} \right) c_{h,t+1}^{-\gamma} \right\}, \tag{2.3}$$

which governs the intertemporal substitution decision of households. Both lower real interest rates and higher future consumption taxes increase the intertemporal policy wedge, $\frac{1}{1+r_{t+1}} \frac{1+\tau_{t+1}^C}{1+\tau_t^C}$, thereby incentivizing households to consume more today.

The labor-leisure equation is given by:

$$l_{h,t}^{\psi} = c_{h,t}^{-\gamma} z_{h,t} w_t \frac{1-\tau_t^L}{1+\tau_t^C}. \tag{2.4}$$

Consumption taxes and labor taxes directly influence the labor supply of households through the intratemporal policy wedge, $\frac{1-\tau_t^L}{1+\tau_t^C}$.

Let $c_t(b,z)$, $\ell_t(b,z)$, and $b_{t+1}(b,z)$ denote the policy functions for consumption, labor supply, and savings, respectively, that satisfy equations (2.2), (2.3), and (2.4) given the household's individual state.

[15] Note that our perfect substitutability result also holds if we assume different rules for the distribution of dividends. We will discuss this further in Section 2.3.

2.2.2 Firms

Final good firms produce in a perfectly competitive market using intermediate goods as inputs. Their decision problem is:

$$\max_{y_{j,t}} \left\{ P_t Y_t - \int_0^1 p_{j,t} y_{j,t} dj \right\},$$

subject to a CES production technology:

$$Y_t = \left(\int_0^1 y_{j,t}^{1/\mu} dj \right)^{\mu},$$

where $y_{j,t}$ denotes the intermediate good produced by firm j and $p_{j,t}$ is the corresponding price. Y_t denotes the final consumption good, P_t denotes the overall price index, and μ determines the degree of substitution among input factors. The aggregate price index is given by:

$$P_t = \left(\int_0^1 p_{j,t}^{1/(1-\mu)} dj \right)^{1-\mu}.$$

Solving the maximization problem yields the demand function of final good firms for the intermediate good j:

$$y_{j,t} = \left(\frac{p_{j,t}}{P_t} \right)^{\frac{\mu}{1-\mu}} Y_t. \tag{2.5}$$

Intermediate goods are produced by a continuum of intermediate good firms in monopolistically competitive markets according to:

$$y_{j,t} = n_{j,t}.$$

Following Correia et al. (2013), we assume that price setting takes place before consumption taxes. As in Calvo (1983), we allow an intermediate good firm to reset its price only with a certain probability, θ. If a firm is allowed to reset its prices, it solves the following dynamic maximization problem:

$$\max_{p_t^*, \{y_{j,s}, n_{j,s}\}_{s=t}^{\infty}} \sum_{s=t}^{\infty} \beta^{s-t} (1-\theta)^{s-t} \left(\frac{p_t^*}{P_s} y_{j,s} - w_s n_{j,s} \right),$$

subject to the final good firms' demand given in (2.5). The optimal price ratio p_t^*/P_t that solves this problem is given by:

$$\frac{p_t^*}{P_t} = \frac{\sum_{s=t}^{\infty} \beta^{s-t} (1-\theta)^{s-t} \left(\frac{P_t}{P_s} \right)^{\frac{\mu}{1-\mu}} Y_s \mu w_s}{\sum_{s=t}^{\infty} \beta^{s-t} (1-\theta)^{s-t} \left(\frac{P_t}{P_s} \right)^{\frac{\mu}{1-\mu}} Y_s}. \tag{2.6}$$

For future reference, let $(1+\pi_{t+1}) = \frac{P_{t+1}}{P_t}$ denote the gross inflation rate.

2.2.3 Policy

We close the model by specifying monetary and fiscal policy.

Monetary policy. To simplify our analytical results in Section 2.3, we assume for now that monetary policy directly controls the real interest rate, r_t. Importantly, our perfect substitutability result does not depend on this simplification and still holds if we assume that monetary policy controls the nominal interest rate. Note that given our timing notation, monetary policy sets r_{t+1} in period t such that r_t is predetermined in period t.

Fiscal policy. The government has expenditures for a fixed amount of government consumption, $\bar{G}$, lump-sum transfers, Tr_t, and for repaying debt, B_t. It finances its expenditures by collecting total tax payments, T_t, and by issuing future debt. The government's budget constraint is given by:

$$\bar{G} + Tr_t + (1+r_t)B_t = B_{t+1} + T_t. \tag{2.7}$$

Total tax payments are given by:

$$T_t = \tau_t^C C_t + \tau_t^L w_t L_t, \tag{2.8}$$

where C_t and L_t denote aggregate consumption and aggregate labor, respectively. For simplicity, we assume for now $\bar{G} = 0$ but relax this assumption in Section 2.4.

2.2.4 Aggregation, Market Clearing, and Equilibrium

The aggregate production function of the economy is given by:

$$S_t Y_t = \int_0^1 n_{j,t} dj \equiv N_t, \tag{2.9}$$

where N_t denotes the aggregate labor demand of the intermediate good firms. S_t measures the efficiency loss that occurs whenever prices differ and is given by:

$$S_t \equiv \int_0^1 \left(\frac{p_{j,t}}{P_t}\right)^{\frac{\mu}{1-\mu}} dj \geq 1.$$

It evolves according to:

$$S_{t+1} = (1-\theta)S_t(1+\pi_{t+1})^{\frac{-\mu}{1-\mu}} + \theta\left(\frac{p^*_{t+1}}{P_{t+1}}\right)^{\frac{\mu}{1-\mu}}. \tag{2.10}$$

Inflation is a function of the optimal relative price of the updating firms:

$$1+\pi_t = \left(\frac{1-\theta}{1-\theta(\frac{p_t^*}{P_t})^{\frac{1}{1-\mu}}} \right)^{1-\mu}. \tag{2.11}$$

The distribution of households over their individual states, $\Gamma_{t+1}(\mathcal{B}, z')$, evolves following the exogenous Markov chain for the productivity level and the endogenously derived savings policy functions of the households. Formally:

$$\Gamma_{t+1}(\mathcal{B}, z') = \int_{\{(b,z):b_{t+1}(b,z)\in\mathcal{B}\}} Pr(z'|z)d\Gamma_t(b,z) \tag{2.12}$$

for all sets $\mathcal{B} \subset \mathbb{R}$. Aggregate labor supply, consumption, and savings are:

$$L_t = \int_0^1 z\ell_t(b,z)d\Gamma_t(b,z), \tag{2.13}$$

$$C_t = \int_0^1 c_t(b,z)d\Gamma_t(b,z), \tag{2.14}$$

and

$$B_{t+1}^d = \int_0^1 b_{t+1}(b,z)d\Gamma_t(b,z), \tag{2.15}$$

respectively.

Labor market clearing requires:

$$L_t = N_t, \tag{2.16}$$

the bond market clears when:

$$B_t = B_t^d, \tag{2.17}$$

and the goods market clears when:

$$Y_t = C_t + \bar{G}. \tag{2.18}$$

Dividend payments are given by:

$$D_t = Y_t - w_t N_t. \tag{2.19}$$

Equilibrium. We define an equilibrium of the economy to consist of:

1. Policy and value functions $\{b_{t+1}(b,z),\, \ell_t(b,z),\, c_t(b,z),\, V_t(b,z)\}_{t=0}^{\infty}$ that solve the households' problems,

2. distributions $\{\Gamma_t(b,z)\}_{t=0}^{\infty}$ that evolve according to (2.12),

3. sequences of the aggregate variables

$$X \equiv \left\{C_t, L_t, N_t, Y_t, d_t, i_t, w_t, \pi_t, r_t, p_t^*/P_t, S_t, Tr_t, T_t, \tau_t^C, \tau_t^L, B_t^d, B_t\right\}_{t=0}^{\infty}$$

that satisfy the equilibrium conditions (2.6), (2.7), (2.8), (2.9), (2.10), (2.11), (2.16), (3.21), (2.19), the household aggregation equations (2.13), (2.14), (2.15), as well as the paths for the real interest rate, consumption taxes, labor taxes, and the government debt level to be specified below.

2.3 HANK-UFP

In this section, we prove that HANK-UFP is a perfect substitute for monetary policy in HANK. In particular, we derive a set of sufficient conditions for three *aggregate* fiscal instruments which jointly replicate the consumption and labor supply of *each* household and, thus, the allocation associated with any given change in interest rates.

2.3.1 Perfect Substitutability with Monetary Policy

Monetary policy. Assume a standard perfect foresight monetary policy experiment. The economy is in steady state when in period $t = 0$, monetary policy announces a new path of real rates, $\{r_t^{MP}\}_{t=1}^{\infty}$, with $r_t^{MP} = \bar{r}$ for all $t > s$ for some s. We denote variables associated with this monetary policy experiment with a superscript MP. Consumption taxes, labor taxes, and the government debt level are fixed at their steady state values, that is, for all t, $\tau_t^{L,MP} = \bar{\tau}^L$, $\tau_t^{C,MP} = \bar{\tau}^C$, $B_t^{MP} = \bar{B}$, while transfers, Tr_t^{MP}, adjust to keep the government budget balanced. We focus on the equilibrium in which the economy converges back to steady state for $t \to \infty$.

Monetary policy affects the economy through changing the optimization problem of households. In particular, it changes the households' problem in two ways: first, monetary policy changes the *intertemporal policy wedge*, $\frac{1}{1+r_{t+1}} \frac{1+\tau_{t+1}^C}{1+\tau_t^C}$, in the Euler equation of households (equation (2.3)) which incentivizes households to intertemporally reallocate consumption. Second, monetary policy has effects on the budget constraints of households. Importantly, these effects are not the same across households since, on the one hand, households differ in the composition of their income and, on the other hand, monetary policy has different effects on the various income components of households. We come back to this in the next paragraph.

HANK-UFP. We now describe how fiscal policy replicates the allocation associated with the monetary policy experiment. Assume that the real interest rate is kept at its steady state level, $r_t^{UFP} = \bar{r}\ \forall t$, and fiscal policy changes the paths for its aggregate instruments,

$\tau_t^{C,UFP}$, $\tau_t^{L,UFP}$, and B_{t+1}^{UFP}. The following proposition states our perfect substitutability result between fiscal policy and monetary policy in HANK.

Proposition 1 *Consider HANK-UFP, a fiscal policy scheme which sets the paths for consumption taxes, $\tau_t^{C,UFP}$, labor taxes, $\tau_t^{L,UFP}$, and the government debt level, B_{t+1}^{UFP}, according to the following conditions*

$$\Big(1+\bar{r}\Big)\frac{1+\tau_t^{C,UFP}}{1+\tau_{t+1}^{C,UFP}} = 1 + r_{t+1}^{MP}, \textit{with } \tau_0^{C,UFP} = \bar{\tau}^C, \tag{2.20}$$

$$\frac{1-\tau_t^{L,UFP}}{1+\tau_t^{C,UFP}} = \frac{1-\bar{\tau}^L}{1+\bar{\tau}^C}, \tag{2.21}$$

$$B_{t+1}^{UFP} = \frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C} B_{t+1}^{MP} \tag{2.22}$$

while lump-sum transfers, Tr_t^{UFP}, adjust to keep the government budget constraint balanced. HANK-UFP yields the same allocation as the monetary policy experiment. That is, consumption and labor supply are the same for each household in every period, i.e., $(c_{h,t}^{UFP}, l_{h,t}^{UFP}) = (c_{h,t}^{MP}, l_{h,t}^{MP})$ $\forall h$ and $\forall t$. Hence, conditions (2.20) - (2.22) are sufficient conditions for HANK-UFP to be a perfect substitute for monetary policy.

While we relegate the formal proof of Proposition 1 to Appendix B.1, we now explain the rationale behind our perfect substitutability result. As in Correia et al. (2008) and Correia et al. (2013), HANK-UFP uses consumption taxes and labor taxes to replicate the effects of monetary policy on the policy wedges in the first-order conditions of households. According to condition (2.20), consumption taxes are set such that the intertemporal policy wedge in the Euler equation of each household is the same as in the monetary policy experiment. Intuitively, by changing the ratio of future over current consumption taxes, fiscal policy changes the relative price of current consumption versus future consumption. This way, fiscal policy triggers the same incentive to intertemporally reallocate consumption as a change in the real interest rate. While there are two possibilities to achieve equivalence in the intertemporal policy wedge with monetary policy, namely a pre-announced change in future consumption taxes *or* surprise changes in today's consumption taxes, the second part of condition (2.20) rules out surprise changes as these are not consistent with equivalence in the budget constraint as we will explain below.

Unlike a change in the real interest rate, adjusting consumption taxes changes the *intratemporal policy wedge*, $\frac{1-\tau_t^L}{1+\tau_t^C}$, in the labor-leisure equations of households (equation (2.4)). When labor taxes are set according to condition (2.21), they offset this effect on the labor supply of households.

Furthermore, HANK-UFP ensures that the effects on the budget constraint of each household are the same as with monetary policy. To this end, HANK-UFP replicates the effects of monetary policy on *each component* of households' incomes. Equivalence in the intratemporal policy wedge (condition (2.21)) ensures that households' net wage is the same in consumption value terms as in the monetary policy experiment. Condition (2.20) ensures that the real return on assets in consumption value terms is the same as in the monetary policy experiment. In particular, a pre-announced change in future consumption taxes changes $(1+\bar{r})\frac{b_{h,t+1}^{UFP}}{1+\tau_{t+1}^{C,UFP}}$ but it leaves $(1+\bar{r})\frac{b_{h,t}^{UFP}}{1+\tau_{t}^{C,UFP}}$ unchanged just like an interest rate change changes the consumption value of households' assets in the next period, $(1+r_{t+1}^{MP})\frac{b_{h,t+1}^{MP}}{1+\bar{\tau}^{C}}$, by changing its return, $1+r_{t+1}^{MP}$, but leaves the consumption value of households' assets in this period, $(1+r_{t}^{MP})\frac{b_{h,t}^{MP}}{1+\bar{\tau}^{C}}$, unchanged. As mentioned above, this is not feasible with a surprise change in consumption taxes today as this would already affect $(1+\bar{r})\frac{b_{h,t}^{UFP}}{1+\tau_{t}^{C,UFP}}$.

Equivalence in the real return on assets implies that households want to save the same amount in consumption value terms as in the monetary policy experiment, that is $b_{h,t+1}^{UFP} = \frac{1+\tau_{t}^{C,UFP}}{1+\bar{\tau}^{C}} b_{h,t+1}^{MP}$. Such savings positions also ensure that each household's asset income is the same as with monetary policy. However, to render it feasible for each household to increase her asset position by the factor $\frac{1+\tau_{t}^{C,UFP}}{1+\bar{\tau}^{C}}$, aggregate asset supply also needs to change by this factor which is the case given that debt dynamics follow condition (2.22). At the same time, debt policies following condition (2.22) shift the same resources measured in consumption value terms to the government as monetary policy. Conditions (2.20)-(2.22) together with the government budget constraint (2.7) then yield the following transfer path:

$$Tr_{t}^{UFP} = \frac{1+\tau_{t}^{C,UFP}}{1+\bar{\tau}^{C}} Tr_{t}^{MP} + D_{t}\left(\frac{1+\tau_{t}^{C,UFP}}{1+\bar{\tau}^{C}} - 1\right). \tag{2.23}$$

Hence, transfers follow the path of transfers associated with the monetary policy experiment adjusted for the change in consumption value. In addition, transfers compensate for the change in consumption value of the dividend income.[16] This reflects that tax revenues are different with HANK-UFP and monetary policy as the change in consumption taxes applies to total output, Y_t, whereas the change in labor taxes only applies to total labor income, $w_t N_t = Y_t - D_t$. Overall, equation (2.23) implies that also the lump-sum income component of households is the same in consumption value terms as with monetary policy.

In sum, HANK-UFP replicates the effects of monetary policy both on the policy wedges in the first-order conditions and on the budget constraints of households. Thus, each household faces the same optimization problem with HANK-UFP and monetary policy which implies that both policies induce the same consumption and labor supply of each household. As neither the interest rate nor fiscal policy variables are part of the firms' equilibrium equations, equivalence in each household's behavior generates the same allocation through general equilibrium.

[16] Note that the same compensation for the loss in consumption value of dividends can be achieved if we allow for a lump-sum transfer to firms. In this case, the aggregate transfer payment to firms would equal the second term in equation (2.23). Hence, our perfect substitutability result does not depend on the assumption of lump-sum dividends but can also be achieved with any other assumption on how dividends are distributed.

Policy-induced redistribution. A corollary of our perfect-substitutability result is that monetary policy and HANK-UFP induce the same redistribution among households. We now show that HANK-UFP achieves this by replicating the redistribution of monetary policy through the policy block, that is, the partial equilibrium redistribution through changes in monetary and fiscal policy variables. To capture this policy-induced redistribution formally, we define the policy-exposure of each household by $\Xi_{h,t}$. This captures the partial-equilibrium changes of resources for household h which are induced by changes in policy variables assuming fixed households' and firms' behavior, that is, $(c_{h,t}, l_{h,t}, \frac{b_{h,t+1}}{1+\tau_t^C}, d_t, w_t) = (\bar{c}_h, \bar{l}_h, \frac{\bar{b}'_h}{1+\bar{\tau}^C}, \bar{d}, \bar{w})$.[17]

Lemma 1 *The policy-exposure of household h towards monetary policy is given by* $\Xi_{h,t}^{MP} = \bar{B}(\bar{r} - r_t^{MP}) - \bar{b}_h(\bar{r} - r_t^{MP})$*. Consider periods of expansionary monetary policy, that is* $r_t^{MP} < \bar{r}$*. Then* $\Xi_{h,t}^{MP} < 0$ *for* $\bar{b}_h > \bar{B}$ *and* $\Xi_{h,t}^{MP} > 0$ *for* $\bar{b}_h < \bar{B}$*. That is, expansionary monetary policy redistributes from households that have a higher asset position than the average to households that have a lower asset position than the average. The opposite is true for periods with contractionary monetary policy. With HANK-UFP, the households' policy exposure is given by* $\Xi_{h,t}^{UFP} = \frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C}\Xi_{h,t}^{MP}$*. Hence, HANK-UFP replicates the policy exposure of monetary policy for each household in consumption value terms.*

Again, we relegate the proof of Lemma 1 to Appendix B.2 and focus here on the intuition. A decrease in interest rates generates a negative wealth effect on assets from which households suffer in proportion to their asset holdings. At the same time, lower interest rates shift resources to the government as the government issues the assets and now has lower interest payments. Hence, lower interest rates imply a redistribution from asset holders to the government. As this implies an increase in lump-sum transfers, expansionary monetary policy redistributes from asset-rich to asset-poor households.[18] HANK-UFP induces the same negative wealth effect on assets from which households suffer in proportion to their asset holdings—exactly as with monetary policy. The reason is that higher consumption taxes decrease the consumption value of assets thereby inflating away the buffer stock of households. As a consequence, households buy additional government debt in proportion to their asset holdings. Since this expansion in debt increases lump-sum transfers to all households, the redistribution through changes in policy variables is the same as with monetary policy.

Relation to perfect substitutability in RANK. There are three differences to the perfect substitutability result in RANK by Correia et al. (2013). First, consumption taxes

[17]The bar on top of aggregate variables denotes their respective steady state values. The bar on top of choice variables of households denotes household h's behavior in stationary equilibrium. As the savings of a household determine its individual state in the next period, we need to adjust the savings for the consumption value. Otherwise, this partial equilibrium decomposition would compare different households in the monetary policy experiment and in the HANK-UFP case as discussed above.

[18]In Section 2.3.2, we show that our perfect substitutability result does not depend on the assumption that transfers adjust after a monetary policy shock but also holds when we assume any other fiscal response.

and labor taxes set according to conditions (2.20) and (2.21), respectively, are sufficient for perfect substitutability with monetary policy in RANK. Our analysis shows that these two instruments are no longer sufficient in HANK. With only tax policies, fiscal policy replicates the effects of monetary policy on the policy wedges in the first-order conditions of households but not on their budget constraints such that households are differently exposed to these tax policies and monetary policy. In HANK, these heterogeneous effects on households' budget constraints matter. This is because first, the heterogeneous exposure of households prevents cross-sectional equivalence with monetary policy. Second, it also breaks the aggregate equivalence since households are heterogeneous in their MPCs. In addition, with only tax policies, the value of total assets in consumption value terms changes which has real effects since households hold these assets for self-insurance purposes. In our numerical analyses in Section 2.4.4, we quantify the shortcomings of this fiscal policy scheme and show them to be substantial.

Second, in RANK, it does not matter whether equivalence in the intratemporal policy wedge is achieved via a surprise change in today's consumption taxes or a pre-announced change in future consumption taxes as both alternatives are consistent with perfect substitutability with monetary policy. As explained above, these two alternatives have different dynamic effects on the asset income of households, but with a representative, permanent-income consumer these differences do not affect her behavior as long as the effects on the policy wedges in the first-order conditions are the same. In contrast, in HANK, only pre-announced changes of future consumption taxes are consistent with perfect substitutability as only these are consistent with equivalence in the sequence of budget constraints of each household.

Third, in RANK, unconstrained monetary policy implements the first-best allocation if lump-sum transfers are available and, consequently, UFP replicates this first-best allocation. Unlike in RANK, unconstrained monetary policy does not necessarily implement the first-best allocation in our HANK framework and, thus, neither does HANK-UFP. Yet, Proposition 1 shows that fiscal policy can always at least implement the welfare associated with unconstrained monetary policy.

2.3.2 Model Extensions

We now first show that our perfect substitutability result does not depend on our assumption on the fiscal response to monetary policy and that second, it also holds when we extend our model by sticky wages and investment.

Alternative fiscal responses to monetary policy

The HANK literature highlights the importance of the fiscal response to monetary policy (see Kaplan et al. (2018) and Auclert et al. (2020)). In particular, Kaplan et al. (2018) distinguish three different fiscal responses: first, transfers adjust to balance the government budget after a change in monetary policy which is our baseline. Second, debt adjusts in the short-run and transfers ensure that government debt returns to its original steady state level in the long-run. Third, government spending adjusts. Our perfect substitutability result between fiscal policy and monetary policy does not depend on the fiscal response we assume. If debt adjusts in the short-run and transfers bring back government debt in the long-run in the monetary policy experiment, the conditions in Proposition 1 are still sufficient. If fiscal policy adjusts government spending in response to monetary policy, Proposition 1 needs to be extended by the following condition for government spending: $G_t^{UFP} = G_t^{MP}$.

Sticky Wages

We now show that Proposition 1 is still sufficient for perfect substitutability between HANK-UFP and monetary policy in a model in which both prices and wages are sticky. We model sticky wages as in Auclert et al. (ming). We here only sketch the modifications of the model and leave the details for Appendix B.3. Instead of each household deciding on her labor supply, labor hours are now determined by the labor demand of a continuum of monopolistically competitive unions. We assume quadratic utility costs of adjusting the nominal wage w_{kt} set by union k, by allowing for an extra additive disutility term $\frac{\nu}{2}\int_k \left(\frac{w_{kt}}{w_{kt-1}} - 1\right)^2 dk$ in households' utility (C.33). In every period t, union k sets a common wage for each of its members, and calls upon its members to supply hours according to a uniform rule, so that $l_{hkt} = L_{kt}$. The union sets w_{kt} to maximize the average utility of its members given this allocation rule. Real wages evolve according to $\frac{1+\pi_t^W}{1+\pi_t} = \frac{w_t}{w_{t-1}}$, where π_t^W is the nominal wage inflation which evolves according to an aggregate non-linear wage Phillips curve

$$\pi_t^w(1+\pi_t^w) = \frac{\epsilon_w}{\nu}\int L_t\left(L_t^{\psi} - \frac{\epsilon_w - 1}{\epsilon_w}\frac{1-\tau_t^L}{1+\tau_t^C} w_t z_{ht} c_{ht}^{-\gamma}\right)dh + \beta\pi_{t+1}^w(1+\pi_{t+1}^w), \tag{2.24}$$

where ϵ_w denotes the elasticity of substitution among unions.

Note that labor supply is now exogenous to households and, thus, every solution to the household problem now only needs to satisfy equations (2.2) and (2.3). Introducing sticky wages does not change either of the two equations compared to our baseline model with flexible wages. Labor supply is now pinned down by the wage Phillips curve (2.24) instead of households' individual labor-leisure equations. However, this Phillips curve is equally affected by HANK-UFP and by monetary policy given that condition (2.21) generates the same intratemporal policy wedge with both policies and given that each household consumes the same and, thus, has the same marginal utility of consumption with both policies. Hence, also with sticky wages, all equilibrium conditions are equally affected by HANK-UFP and by monetary policy and, thus, our perfect substitutability result holds.

Extension to investment

In this section, we show that fiscal policy can also be a perfect substitute for monetary policy if we extend our model by capital. To this end, we assume that intermediate goods firms produce according to a typical Cobb-Douglas production function with labor and capital as inputs (see Appendix B.4 for details). We allow for a linear capital subsidy, τ_t^F, which is directly paid to the intermediate goods firms. The firms' first-order conditions yield the following term for real marginal cost:

$$mc_t = \left(\frac{1}{\alpha}\right)^{\alpha} \left(\frac{1}{1-\alpha}\right)^{1-\alpha} (r_t^{k,F} - \tau_t^F)^{\alpha} (w_t)^{1-\alpha}, \tag{2.25}$$

where α is the share of capital and $r_t^{k,F}$ is the rental rate of capital. An increase in capital subsidies increases the demand for capital in the same way as expansionary monetary policy since both policies equally affect the policy wedge in the firms' first-order conditions. Thus, HANK-UFP replicates the incentives to use capital induced by monetary policy by setting the capital subsidy according to the following condition (assuming $\tau_t^{F,MP} = \bar{\tau}^F = 0$ without loss of generality):

$$r_t^{k,F,MP} = r_t^{k,F,UFP} - \tau_t^{F,UFP}. \tag{2.26}$$

The economy's capital stock evolves according to $K_{t+1} = I_t + (1-\delta)K_t$, where δ is the depreciation rate of capital and the investment good, I_t, is produced with the same technology as the final consumption good. We assume that the depreciation of capital is replaced through maintenance and, thus, $r_t^K = r_t^{K,F} - \delta$ is the net return on capital. The capital stock is held by the households and, thus, monetary policy also affects households' budget constraints through changing the return on capital. The effects on households' budget constraints and, thus, the specification of HANK-UFP depend on the degree of substitutability between bonds and capital from the perspective of households. We next study a model in which capital is a perfect substitute to bonds and one in which capital is illiquid and, thus, an imperfect substitute to bonds.

Bonds and capital as perfect substitutes. We start by specifying the conditions for HANK-UFP when capital is a perfect substitute for bonds. Perfect substitutability of both assets implies that the return on both assets is the same and, hence, $r_t = r_t^k$. In this case, the budget constraint of household h is given by:

$$c_{h,t} = (1+r_t)\frac{b_{h,t} + k_{h,t}}{1+\tau_t^C} - \frac{b_{h,t+1} + k_{h,t+1}}{1+\tau_t^C} + \frac{1-\tau_t^L}{1+\tau_t^C} w_t z_{h,t} l_{h,t} + \frac{D_t + Tr_t}{1+\tau_t^C} \tag{2.27}$$

$$b_{h,t+1}, k_{h,t+1} \geq 0,$$

where $k_{h,t}$ is the amount of capital held by household h.[19]

Consider the following modified condition for the government debt level:

$$B_{t+1}^{UFP} = \frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C}\bar{B} + \left(\frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C} - 1\right)K_{t+1}^{MP}. \tag{2.28}$$

HANK-UFP consisting of conditions (2.20), (2.21), (2.28), and (2.26) is a perfect substitute for monetary policy. We relegate the proof to Appendix B.4.5 and focus here on the intuition of this result.

Equivalence in the intertemporal policy wedge in the Euler equations of households implies that each household wants to save the same amount in consumption value terms:

$$b_{h,t+1}^{UFP} + k_{h,t+1}^{UFP} = \left(\frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C}\right)(b_{h,t+1}^{MP} + k_{h,t+1}^{MP}). \tag{2.29}$$

When the savings of each household evolve according to (2.29), the budget constraint of each household is the same with HANK-UFP and with monetary policy. Setting the government debt level according to condition (2.28) implies that $K_t^{UFP} = K_t^{MP}$ and that equation (2.29) is feasible for each household.

Hence, in addition to replicating the incentives for firms to use capital, HANK-UFP now also replicates the redistribution from capital holders to firms induced by monetary policy: households compensate for the loss in consumption value of their capital holdings by buying more bonds in proportion to their capital holdings. These additional bonds, in turn, finance the capital subsidies to firms.

Bonds and capital as imperfect substitutes. We now assume that capital holdings are subject to an illiquidity friction which renders capital an imperfect substitute for bonds (as for example in Kaplan et al. (2018) and Bayer et al. (2019)). In this case, $r_t^k = r_t + \epsilon_t^k$, where ϵ_t^k is an endogenous and time-varying spread on illiquid capital that households demand because capital is less liquid and, thus, less suited for self-insurance purposes. As Bayer et al. (2023) show, ϵ_t^k also depends on the amount of bonds. As a consequence, HANK-UFP cannot compensate for the loss in consumption value of savings in capital by increasing the bond supply but it needs to use an additional instrument. More specifically, HANK-UFP needs to issue an additional asset which we label *synthetic capital*, Ω, which has the same pay-off structure and the same illiquidity friction as capital. In other words, synthetic capital is a perfect substitute for capital from the households' perspective but it is not used in production. Accordingly, the government budget constraint is given by:

$$Tr_t + \tau_t^F K_t + (1+r_t^k)\Omega_t + (1+r_t)B_t = B_{t+1} + \Omega_{t+1} + T_t, \tag{2.30}$$

[19]Given $r_t^k = r_t$, the portfolio of each household is indeterminate. For simplicity, we assume that all households have the same bond to capital ratio of $\bar{B}/\bar{K}$ but our results do not depend on this assumption.

and without loss of generality, we assume that $\Omega_t^{MP} = \bar{\Omega} = 0$.

The budget constraint of household h is then given by:

$$\begin{aligned} c_{h,t} = \frac{1+r_t}{1+\tau_t^C} b_{h,t} - \frac{b_{h,t+1}}{1+\tau_t^C} + \frac{1+r_t^K}{1+\tau_t^C}(k_{h,t}+\omega_{h,t}) - \frac{(k_{h,t+1}+\omega_{h,t+1})}{1+\tau_t^C} \\ + \frac{1-\tau_t^L}{1+\tau_t^C} w_t z_{h,t} l_{h,t} + \frac{D_t + Tr_t}{1+\tau_t^C}, \\ b_{h,t+1}, k_{h,t+1} \geq 0, \end{aligned} \tag{2.31}$$

where $\omega_{h,t+1}$ is the amount of synthetic capital household h holds.

Consider the following condition for the issuance of synthetic capital:

$$\Omega_{t+1}^{UFP} = \left(\frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C} - 1 \right) K_{t+1}^{MP}. \tag{2.32}$$

HANK-UFP consisting of conditions (2.20), (2.21), (2.22), (2.26), and (2.32) is a perfect substitute for monetary policy. We relegate the formal proof to Appendix B.4.6 and focus here on the intuition.

Equivalence with monetary policy in the real return on illiquid assets implies that each household wants to save the same amount of illiquid assets in consumption value terms:

$$\omega_{h,t+1}^{UFP} + k_{h,t+1}^{UFP} = \left(\frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C} \right) (\omega_{h,t+1}^{MP} + k_{h,t+1}^{MP}). \tag{2.33}$$

Given that households' portfolios evolve according to $\left(b_{h,t+1}^{UFP}, k_{h,t+1}^{UFP} + \omega_{h,t+1}^{UFP}\right) = \left(\frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C} b_{h,t+1}^{MP}, \frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C}(k_{h,t+1}^{MP} + \omega_{h,t+1}^{MP})\right)$, the budget constraint of each household is equally affected by HANK-UFP and monetary policy. Setting synthetic capital according to condition (2.32) implies that $K_t^{UFP} = K_t^{MP}$ and that equation (2.33) is feasible for each household. This way, HANK-UFP also replicates the redistribution from capital holders to firms: households hold the synthetic capital in proportion to their capital holdings, which, in turn, finances the capital subsidy.

2.4 Circumventing the ELB Constraint

Our perfect substitutability result between fiscal policy and monetary policy is especially relevant when monetary policy is constrained. To illustrate this, we now show how HANK-UFP circumvents the ELB constraint—a natural example of constrained monetary policy. In this section, we first demonstrate that HANK-UFP achieves the same allocation as hypothetically unconstrained monetary policy when a discount factor shock pushes the economy to the ELB. To this end, we now study a numerical exercise of our model in Section 2.2 including sticky wages as described in Appendix B.3. We then show that if the

UFP measure does not include debt policies but only consists of tax policies, output and inflation fall by more than with unconstrained monetary policy while the ELB binds and the economy converges to a new steady state after the ELB stops binding.

Monetary policy. We now assume that monetary policy controls the nominal interest rate and follows a Taylor rule. The central bank sets the nominal interest rate between period t and $t+1$, i_{t+1}, according to:

$$1 + i_{t+1} = \max \left\{ \underline{I}, (1+\bar{i}) \left(\frac{\pi_t}{\bar{\pi}} \right)^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}} \right)^{\phi_Y} \right\}. \tag{2.34}$$

The parameters ϕ_π and ϕ_Y measure how responsive the central bank reacts to deviations in inflation and output, respectively, from steady state. In the case of constrained monetary policy, the Taylor rule is truncated by the ELB. Thus, $\underline{I} = 1$ and nominal interest rates cannot go below zero. In the counterfactual of unconstrained monetary policy, $\underline{I} \longrightarrow -\infty$, and monetary policy follows the Taylor rule without any constraints.

The nominal and the real interest rate are linked via the Fisher equation:

$$1 + r_{t+1} = \frac{1 + i_{t+1}}{1 + \pi_{t+1}}. \tag{2.35}$$

2.4.1 Calibration

Table 2.1 summarizes our calibration which are standard values in the literature. We set

Table 2.1: Calibration of the model.

Parameter	Description	Value
β	Discount factor	0.982
γ	Risk aversion	2
ψ	Inverse of Frisch elasticity	2
μ	Markup	1.2
θ	Price reversion rate	0.15
ϵ^W	Elasticity of substitution among unions	11
ν	Disutility of resetting wage	100
ρ_z	Autocorrelation of idiosyncratic risk	0.966
σ_z	Unconditional variance of idiosyncratic risk	0.501
$\bar{\tau}^C$	Consumption tax rate	5%
$\bar{\tau}^L$	Labor tax rate	28%
$\bar{G}/\bar{Y}$	Government consumption share	0.2
$\bar{Tr}/\bar{Y}$	Transfer share	0.055
$\bar{B}/(4 * \bar{Y})$	Government debt share	0.9
ϕ_π	Inflation Taylor weight	1.5
ϕ_Y	Output Taylor weight	0.125

the households' discount factor, β, such that the annual steady state real interest rate, $\bar{r}$, is 2%. We set both the coefficient for risk aversion, γ, and for the inverse Frisch elasticity, ψ, to 2. The latter reflects the finding of Chetty (2012). Following Christiano et al. (2011), we set the markup parameter, μ, to 1.2, and the price reversion rate, θ, to 0.15. For the degree of wage stickiness, we follow Auclert et al. (ming). In particular, we set the elasticity of substitution among union-specific tasks, ϵ^W, to 11 and the parameter governing the disutility of resetting wages, ν, to 100. If we linearized the wage Phillips curve, these values would imply a slope of 0.1 as in Auclert et al. (ming).

The calibration of the idiosyncratic income risk follows McKay et al. (2016). We assume that households cannot borrow. We choose the labor income risk to approximate the findings of Floden and Lindé (2001). We discretize a quarterly AR(1) process with an autoregressive coefficient of 0.966 and an innovation variance of 0.017 into a Markov chain by using Rouwenhorst (1995)'s method.[20] The resulting Markov chain matches the unconditional and the conditional mean, the unconditional and the conditional variance, and the first-order autocorrelation of the underlying quarterly AR(1) process. For our discretization, we choose an 11-state Markov chain as for example in Auclert et al. (ming).

Following Correia et al. (2013), we set the consumption tax rate, $\bar{\tau}^C$, to 5% and the labor tax rate, $\bar{\tau}^L$, to 28%. As in Christiano et al. (2011), we set government consumption $\bar{G}/\bar{Y} = 0.2$. We set the government debt level to target a quarterly average MPC of 0.16 as in Kaplan et al. (2018). This results in an annual government debt share, $\bar{B}/(4 * \bar{Y})$, of 90%. Balancing the government budget then requires a steady state transfer share of $\bar{Tr}/\bar{Y} = 0.06$. We set the Taylor coefficient on inflation and output to 1.5 and 0.125, respectively, as it is standard in the literature.

2.4.2 Solution Method

We solve the model using the perfect foresight method proposed in McKay et al. (2016). We compute the transition paths of the economy in response to a discount factor shock. Initially, the economy is in steady state. Without fiscal policy interventions, we assume that the economy returns to its old steady state after 250 periods. With HANK-UFP, we assume that the economy has transitioned to its new steady state after 250 periods.

We guess the paths of the prices and the quantities of the variables specified in Section 2.2.4. We then check whether these prices and quantities are consistent with the definition of an equilibrium in Section 2.2.4 in each period. This implies to solve for the aggregate behavior of households given the guessed prices in each period. We use the endogenous grid point method of Carroll (2006) to solve the individual household problem backwards. We use the non-stochastic simulation algorithm in Young (2010) to simulate the distribution of households forward. When the aggregate behavior of households is not consistent with the guessed quantities, we update the guess for prices and quantities.[21]

[20] Floden and Lindé (2001) estimate the annual log wage process assuming that it follows an AR(1) process resulting in an autoregressive coefficient of 0.961 and an innovation variance of 0.426. The annual AR(1) process is simulated by a quarterly AR(1) process with an autoregressive coefficient of 0.966 and an innovation variance of 0.017.

[21] We use an auxiliary model to update our guesses. It approximates the aggregate behavior of households with an auxiliary Euler equation and an auxiliary aggregate wage Phillips curve which contain time-varying

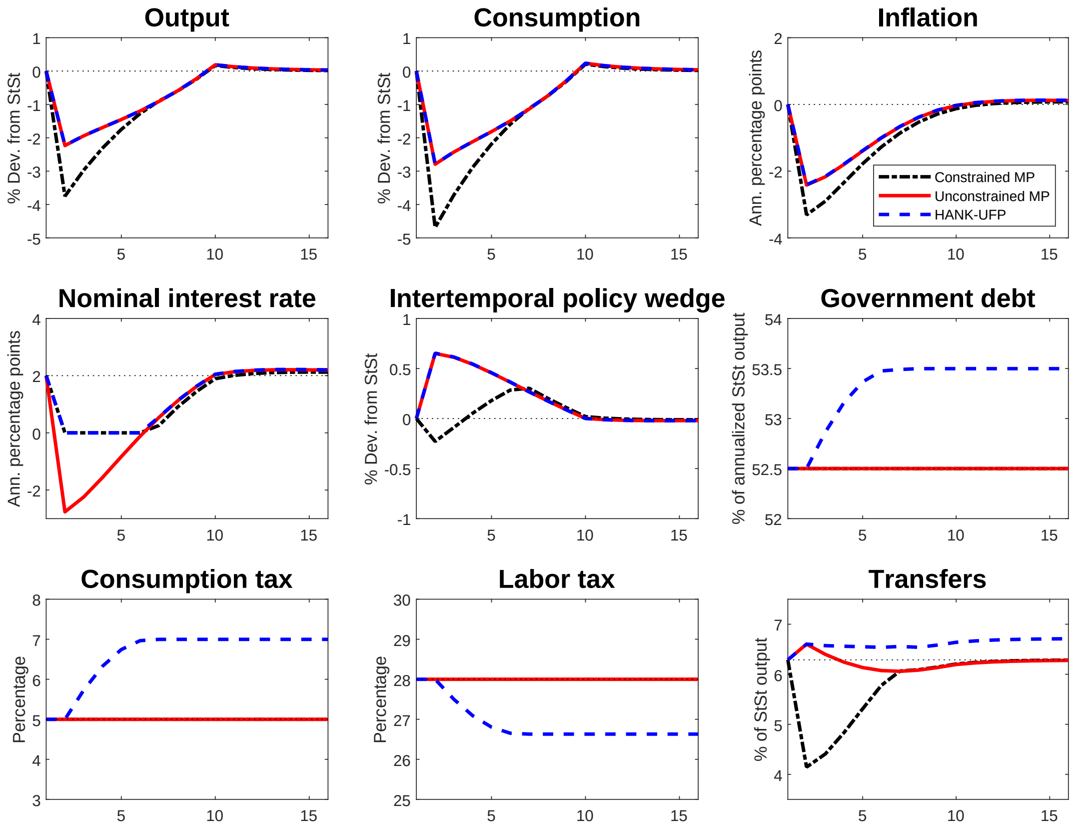

Figure 2.1: Impulse response functions after a shock to the discount factor with a Taylor rule truncated by the ELB ("Constrained MP"), with a Taylor rule without a lower bound ("Unconstrained MP"), and with a truncated Taylor rule and an additional HANK-UFP stimulus ("HANK-UFP"). Horizontal axes denote quarters.

2.4.3 HANK-UFP at the ELB

We follow Christiano et al. (2011) and approximate the effects of a binding ELB by engineering an unexpected temporary increase in the discount factor of households. The discount factor increases by 1.45% for 8 quarters before it jumps back to its steady state level. This brings the economy to the ELB which then binds for 5 quarters. The black, dash-dot lines in Figure 2.1 show the dynamics of macroeconomic aggregates in the *constrained monetary policy case*. In this case, we assume that monetary policy is constrained by the ELB and that there is no fiscal stimulus in the sense that taxes and the government debt level stay constant. Output falls by 3.9%, consumption by 4.9%, and inflation by 3.5 annual percentage points.

How would macroeconomic aggregates react if monetary policy was not constrained by the ELB? The red, solid lines show this *unconstrained monetary policy case*. Without the ELB constraint, the central bank sets negative interest rates for 5 quarters which decrease at most to -2.8 annual percentage points. Output now only falls by 2.2%, consumption by 2.8%, and inflation by 2.4 annual percentage points.

heterogeneity wedges. We solve the auxiliary model with a version of Newton's method and iterate until the aggregate behavior of households is consistent with the guessed quantities and prices.

The blue, dashed lines show the *HANK-UFP case*. The impulse response functions (IRFs) of the macroeconomic aggregates reflect that HANK-UFP achieves the same stabilization as unconstrained monetary policy since both responses lie perfectly on top of each other. Fiscal policy sets the paths for consumption taxes, labor taxes, and the government debt level to replicate the effects of hypothetically unconstrained monetary policy on both the first-order conditions and on the budget constraint of each household. According to condition (2.20), consumption taxes increase while the ELB is binding along a pre-announced path, in total from 5.0% to 7.1%. This way, consumption taxes replicate the effects through the intertemporal substitution channel of unconstrained monetary policy which is reflected by the IRFs of the intertemporal policy wedge in both cases. Labor taxes, correspondingly, decrease in total from 28.0% to 26.6% (condition (2.21)). In line with condition (2.22), government debt increases to a higher level of 91.8% (instead of 90.0%) such that households can hold the same amount of assets in consumption value terms as in the unconstrained monetary policy case. As equation (2.23) shows, this implies that transfers follow the path of transfers in the unconstrained monetary policy case but overshoot them to compensate for the loss in consumption value of the lump-sum income component. At most, transfers increase from 5.5% to 6.9% of GDP.

Cross-sectional equivalence. Section 2.3 shows that HANK-UFP replicates monetary policy by replicating the consumption and labor supply of each household in every period. Obviously, this result also holds in our numerical example. Hence, the welfare of each household and the paths for consumption inequality are the same. The equivalence also holds for the paths of wealth and income inequality when adjusting wealth and income for consumption value terms.

Return to steady state. Given repeated ELB crises, the labor tax might eventually become negative. To avoid that the labor tax turns into a labor subsidy, HANK-UFP can be reversed in good times once the constraint on interest rates is relaxed. This requires a decreasing path for consumption taxes, an increasing path for labor taxes, and a decrease in the government debt level accompanied by expansionary monetary policy. In other words, reversing HANK-UFP allows the fiscal authority to bring back the three fiscal instruments to their original steady state level.

2.4.4 Importance of Debt Policies

We now analyze the fiscal policy measure of Correia et al. (2013) —which only consists of tax policies—through the lens of our HANK model. Consumption taxes and labor taxes are set as in Proposition 1 to replicate the effects on the policy wedges in the first-order conditions of households while government debt is not adjusted for the consumption value but the same as in the unconstrained monetary policy case. Figure 2.2 compares this

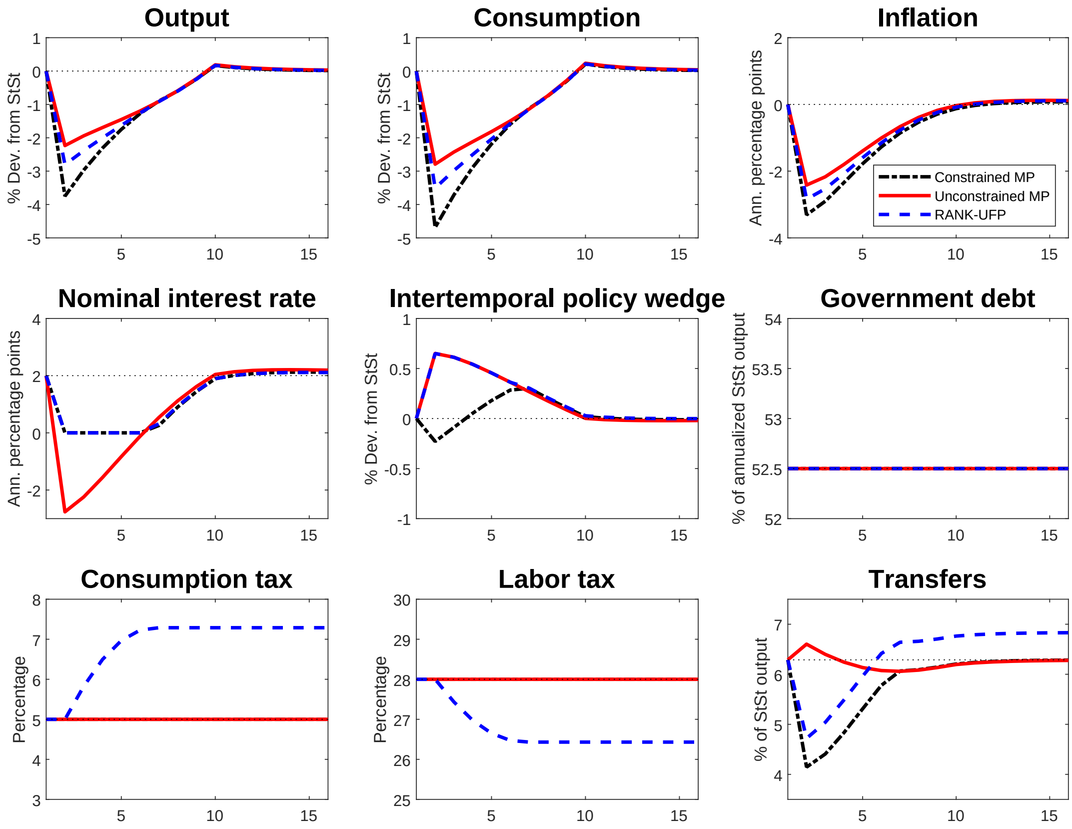

Figure 2.2: Impulse response functions after a shock to the discount factor with a Taylor rule without a lower bound ("Unconstrained MP") and with a truncated Taylor rule and an additional RANK-UFP stimulus ("RANK-UFP"). Horizontal axes denote quarters.

RANK-UFP measure (blue, dashed lines) with unconstrained monetary policy (red, solid lines) in response to the same discount factor shock as in Section 2.4.3. Note that the unconstrained monetary policy case as well as the constrained monetary policy case are the same as in Figure 2.1.

There are two take-aways: first, RANK-UFP does not achieve the same stabilization as unconstrained monetary policy while the ELB is binding as reflected in Figure 2.2. The reason is that lump-sum transfers are lower than with HANK-UFP while the ELB binds since government debt does not increase in these periods. Hence, RANK-UFP does not provide additional resources to high-MPC households while the ELB is binding.[22] Accordingly, output drops on impact by −2.9%, consumption by −3.6%, and inflation by −3.0 annual percentage points.

The second take-away is that in the long-run, the real interest rate does not converge back to its original steady state level of 2% annually, but converges to a lower steady state level of 1.9% annually. The reason is that the supply of assets in consumption value terms is lower than in the unconstrained monetary policy case since government debt does not increase. Hence, households cannot hold the same amount of savings in consumption value terms and, hence, they are worse insured against their idiosyncratic income risk. This increases the inefficiencies from incomplete markets which corresponds to the negative effect of lower asset supply highlighted by Guerrieri and Lorenzoni (2017). As we will show below, this has a sizeable detrimental impact on households' welfare.[23] Our analysis shows that in HANK models, tax policies that interact with the precautionary savings motive of households are not neutral but can have a quantitatively significant impact on the economy in the long-run. Accordingly, debt policies play a crucial role in balancing these effects. This is in stark contrast to RANK models in which there is no precautionary savings motive and the asset demand is perfectly elastic with respect to the real interest rate in the long-run.[24]

Cross-sectional outcomes. As macroeconomic dynamics are not the same, cross-sectional outcomes also differ with RANK-UFP. We summarize the cross-sectional differences by comparing the welfare implications of RANK-UFP and HANK-UFP on each household. We compute the consumption compensation of each household in the economy with RANK-UFP and with HANK-UFP (which is the same as with unconstrained monetary policy).[25]

[22]The stimulative effect of transfer policies is a common feature in HANK models as Ricardian equivalence does not hold (Oh and Reis 2012, Hagedorn et al. 2019, Bayer et al. 2020, Wolf 2021).

[23]In a previous version of the paper, we assumed flexible wages in our numerical example. In that case, the worsened insurance possibilities of households due to RANK-UFP also result in a decrease in the effective labor supply of households and, thus, in a lower steady state output. With sticky wages, this effect is infinitesimal due to the assumption that labor supply is determined by unions.

[24]The convergence to a new steady state can be prevented if RANK-UFP is reversed after the ELB stops binding. This would imply a decreasing path of consumption taxes and an increasing path of labor taxes supported by expansionary monetary policy. Yet, given the non-equivalence of RANK-UFP and monetary policy, this reversal would again not be neutral on the allocation.

[25]We compute the consumption compensation as the consumption increase that is additionally necessary in the baseline of the constrained monetary policy case such that each household is indifferent between the baseline and the two policy cases (RANK-UFP and HANK-UFP). Given our specification of preferences,

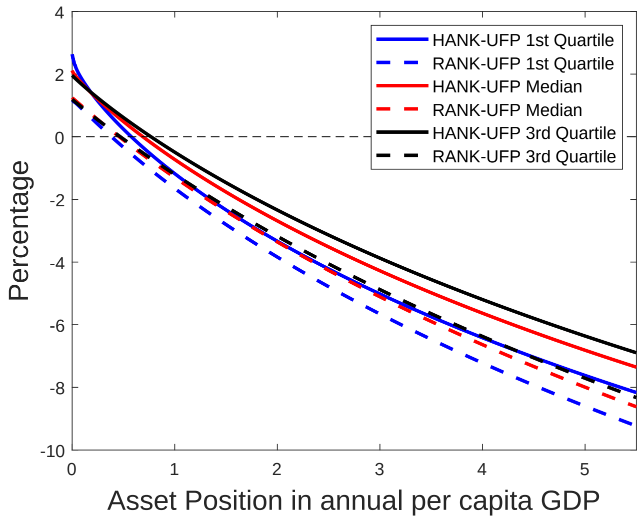

Figure 2.3: Consumption compensation for 4 quarters for each household such that she is indifferent between the respective policy and constrained monetary policy.

Our welfare computation shows that each household is worse-off with RANK-UFP than with HANK-UFP, independent of her idiosyncratic state. Figure 2.3 shows this for a subset of households. It depicts the consumption compensation along the wealth distribution with RANK-UFP and with HANK-UFP for the households whose productivity levels correspond to the 1st quartile (blue), the median (red), and the 3rd quartile (black) of the productivity distribution. The solid lines depict the consumption compensation in the HANK-UFP case and the respective dashed lines depict the consumption compensation in the RANK-UFP case. Figure 2.3 shows that the solid lines always lie above the respective dashed lines indicating the welfare gain of each of these households with HANK-UFP compared to RANK-UFP.[26] Overall, our welfare analysis highlights that adding debt to the fiscal policy mix induces large welfare gains for each household.

2.5 Conclusion

We show that fiscal policy can be a perfect substitute for monetary policy in a HANK model as it can replicate the allocation of monetary policy. The insight is that by changing consumption taxes, labor taxes, and the government debt level, fiscal policy can manipulate the optimization problem of each household in the same way as a change in interest rates:

we cannot compute lifetime consumption compensation. Thus, we compute the consumption compensation for 4 quarters as in Kekre (2021).

[26]Note that all lines of both RANK-UFP and HANK-UFP cross the x-axis at some point. This indicates that high-asset households are worse off with the HANK-UFP stabilization (and, equivalently, with the unconstrained monetary policy stabilization). The reason is that the stabilization policy reduces their wealth significantly in consumption value terms. This effect outweighs their welfare loss out of a recession caused by the ELB in the baseline case of constrained monetary policy. Yet, the average consumption compensation of HANK-UFP (and equivalently unconstrained monetary policy) is 1.42% reflecting that HANK-UFP would be highly beneficial from a Utilitarian social planner perspective compared to RANK-UFP which only has an average consumption compensation of 0.02%.

these tax and debt policies jointly replicate the effects of interest rate changes on the policy wedges in the first-order conditions and the budget constraint of each household. Our perfect substitutability result is especially relevant when monetary policy is constrained—be it due to a binding lower bound, a currency union, or an exchange rate peg—since it implies that fiscal policy can circumvent these constraints. Unlike analyses with a representative agent, our analysis shows that including debt policies in the fiscal policy mix is necessary for perfect substitutability with monetary policy in HANK. Moreover, we highlight at the ELB that not using debt policies has quantitatively important consequences for cross-sectional and aggregate outcomes. When the fiscal authority uses only tax policies, it cannot replicate macroeconomic aggregates in the short-run and induces detrimental effects in the long-run.

We conclude by pointing out two avenues for future research. First, we have shown that a fiscal policy consisting of consumption taxes, labor taxes, and the government debt level can always implement the welfare of unconstrained monetary policy in HANK. Our conjecture is that fiscal policy can generate a higher welfare than unconstrained monetary policy. Second, we prove our perfect substitutability result in a perfect foresight environment. This means it applies to models without aggregate risk or equivalently to the model's first-order perturbation with aggregate risk (see Boppart et al. (2018) and Auclert et al. (2021)). Our conjecture is that our result still holds with higher-order aggregate risk. The reason is that higher-order aggregate risk might change the path of interest rates that the central bank wants to implement. However, the effects of this path of interest rates on the households' optimization problems would still be replicated by our fiscal policy scheme as our fiscal policy scheme replicates the allocation of any path of interest rates. We leave a formal analysis with higher-order aggregate risk for future research.

Chapter 3

A HANK^2 Model of Monetary Unions

with Christian Bayer, Alexander Kriwoluzky, and Gernot J. Müller

Abstract

How does a monetary union alter the impact of business cycle shocks at the household level? We develop a Heterogeneous Agent New Keynesian model of two countries (HANK^2) and show in closed form that a monetary union shifts the adjustment to a shock horizontally across countries, within the brackets of the union-wide wealth distribution, rather than vertically, that is, across the brackets of the union-wide wealth distribution. Calibrating the model to the euro area reveals that a monetary union alters the impact of shocks most strongly in the tails of the wealth distribution but leaves the middle class almost unaffected.

Keywords: OCA theory; two-country model; monetary union; monetary policy; household heterogeneity; inequality
JEL-Classification: F45, E52, D31 This version was published as
Bayer, C., Kriwoluzky, A., Müller G., and F. Seyrich (2024). Unconventional Fiscal Policy in a Heterogeneous-Agent New Keynesian Model *Journal of Monetary Economics* 103579. doi: `https://doi.org/10.1016/j.jmoneco.2024.103579`

3.1 Introduction

Following the seminal work of Mundell (1961), Optimum Currency Area theory analyzes the costs and benefits of monetary unions at the level of regions or countries. Likewise, the policy debate is framed in such terms, as the 20-year plus history of the euro illustrates: discussions of whether specific *countries* would have been better off without the euro abound. Heterogeneity is at the heart of the issue: if countries differ, say, because of country-specific shocks, one (monetary) policy doesn't fit all. However, heterogeneity across households—in terms of income, wealth and shocks—dwarfs the heterogeneity across countries. Hence, we offer a change of perspective. We focus on households rather than countries and ask: How does a monetary union alter the impact of business cycle shocks at the household level?

To answer this question, we propose a Heterogeneous Agent, New Keynesian model of two countries: HANK^2. The model features incomplete markets, idiosyncratic risk, and self-insurance in a standard New Keynesian two-country setup. It is therefore able to capture key features of the business cycle and the wealth distribution; and we may use it to analyze how a monetary union alters—at three levels of aggregation—the effects of country-specific business-cycle shocks relative to a scenario where independent monetary policies are in place. First, we show that when aggregated across both countries and households, macroeconomic dynamics are independent of whether there is a monetary union or not. Second, the overall effect of the shock for households in specific brackets of the wealth (and income) distribution, aggregated across countries of residence, does not depend on whether monetary union is in place or not. In other words, the monetary union itself does not shift the impact of the shock *vertically* across wealth classes. Instead, and this is our third and main result, it shifts the impact of the shock *horizontally* across borders within the brackets of the wealth distribution. Quantitatively, this shift is strongest in the tails of the distribution.

We first study a small-scale version of the model similar to Auclert et al. (2021) where we make a number of simplifying assumptions. These assumptions allow us to obtain strong analytical knife-edge results. First, we restrict the two countries in the model to be symmetric, except for the occurrence of country-specific shocks. What is more, we abstract from capital accumulation and allow for trade in liquid, one-period debt only. Further, we abstract from price stickiness and maintain wage rigidities as the only nominal friction. A key difference to Auclert et al. (2021) is that we consider a two-country model rather than a small-open economy. Relying on a sequence-space representation as in Auclert et al. (ming) or McKay and Wolf (2022), we then show that it is possible to cast the union-wide dynamics of the simplified model into the canonical form which is familiar from the textbook-version of the Representative Agent New Keynesian (RANK) model Galí (2015).

To assess the quantitative relevance of the analytical results, we consider a richer version of the model. Specifically, we introduce capital formation, portfolio choice, and price rigidities building on the medium-scale HANK model in Bayer et al. (ming). We calibrate this version of the model to capture key features of (and asymmetries between) the Italian economy and the German one at the household level. At this level, Italy and Germany are very different, notably in terms of wealth inequality. According to a number of indicators, wealth inequality is significantly higher in Germany. We show that the model is able to account for these structural differences. At the macro level, we maintain the assumption of identical frictions. These have been the subject of earlier research, which we review below. Still, we verify that the model is able to capture key aspects of the business cycle, including its co-movement across Italy and Germany. Importantly, we find that the results that we establish for the simplified model approximately hold for the medium-scale HANK model, even though household-level heterogeneity differs considerably across countries.

In more detail, we present our main results as we put forward three propositions. First, we show that whether countries operate a monetary union or independent monetary policies makes no difference for how country-specific shocks play out at the union level: Monetary union is irrelevant for union-wide dynamics. This result holds exactly for a first-order approximation in the aggregate states and under the assumption that countries are perfectly symmetric (except for the incidence of shocks). Moreover, we require that the monetary union is designed in such a way that that the (implied) monetary policy rule for the union-wide interest rate does not differ from the case of independent monetary policies. This holds if we assume—in line with actual practice in the EA—that the common monetary policy adjusts the policy rate to the average inflation rate (and possibly the output gap) in both countries, to which we refer as "Home" and "Foreign". Against this background the irrelevance result is intuitive. Under a monetary union, monetary policy does not fit all: relative to a benchmark with independent monetary policies, the common policy responds too much in one country and too little in the other. It follows that macro dynamics at the country level do very much depend on whether countries operate a monetary union or not. But when countries are symmetric the changes induced by the monetary union in both countries offset each other such that union-wide dynamics do not change with monetary union.

Our second proposition concerns the household level. Taking a union-wide perspective and aggregating households across countries of residence within the brackets of the wealth distribution, we find that the impact of a business cycle shock for specific wealth classes does not depend on whether there is a monetary union or not. Put differently, just like with union-wide aggregate dynamics, monetary union is also irrelevant for the impact of shocks along the union-wide wealth distribution. It does not, say, shift the adjustment vertically from the rich to the poor or vice versa. Intuitively, how saving and consumption at the household level change in response to shocks within one country does depend on whether there is a monetary union in place or not because these depend on the price adjustments within that country. But given that a monetary union is irrelevant for union-wide price paths, it follows that the changes that a monetary union induces for the response of a generic household in Home are perfectly offset by the changes of its "twin" in Foreign—where a twin is defined in terms of its location in the income and wealth space. Aggregating across countries of residence, we thus find the overall adjustment unchanged within the brackets of the wealth distribution.

Our third proposition is implicit in the argument above. It establishes that monetary union is potentially very relevant for the impact of country-specific shocks along the wealth distribution within a country. More specifically, comparing the outcome under a monetary union to the outcome under independent monetary policies, we observe that the union shifts the impact of shocks horizontally across borders within the brackets of the wealth distribution. Put differently, the impact of shocks changes for specific households at the expense of their twins in the other country: in the face of specific shocks, the poor (rich) in one country benefit from union membership at the expense of the poor (rich) in the other country. Hence, monetary union makes a difference for how shocks impact the rich and the poor within a country.

We simulate the calibrated model and verify that our main results hold approximately once we allow for asymmetries in terms of household-level heterogeneity across countries. In particular, as we study the adjustment to country-specific shocks we find—consistent with Proposition 1—that union-wide aggregate dynamics are basically independent of whether there is a monetary union in place or not. In contrast, country-level dynamics change fundamentally due to the monetary union. In this regard, our model simulation confirms the classic notion that one size doesn't fit all. The response of Gini coefficients in our simulations suggests that Proposition 2 also holds approximately in the asymmetric model.

Last, we perform a quantitative analysis that relates to Proposition 3. We compute the consumption equivalent welfare variation of a shock as a comprehensive (ex-post) measure of its impact and find that a monetary union induces strong changes in this measure in the tails of the wealth distribution, both in Home and Foreign. These changes can be traced back to how a monetary union changes the interest-rate dynamics to which households in the tails of the wealth distribution—rich and poor—are more exposed than the middle class which neither borrows nor saves much (in excess of what it implicitly owes through government debt). We find accordingly, that monetary union does not change the impact of shocks for the middle class. This result offers a fresh perspective on the euro. During its 20-year-plus history, the euro area witnessed various political movements in several of its member states that campaigned against the euro, yet their appeal to the electorate turned out to be limited. Our analysis offers an explanation for why this occurs.

Related literature. Our analysis builds on two earlier generations of OCA theory. The first generation stresses that countries should be sufficiently homogeneous to qualify as an OCA. The original contribution of Mundell (1961) emphasizes that economic regions as opposed to nation-states or countries are the relevant category when it comes to operating a common currency. We thus follow Mundell's lead as we attempt to shift the focus away from countries (and towards households). Other contributions to the first generation of OCA theory stress the role of trade openness and the asymmetry of shocks McKinnon (1963);

Kenen (1981); Bayoumi and Eichengreen (1992); Krugman (1993). Lastly, influential work has emphasized the potential endogeneity of the OCA criteria Frankel and Rose (1998); Rose (2000).

The second generation of OCA theory zooms in on specific aspects, notably on the trade-offs faced by monetary and fiscal policy in monetary unions as well as on the conduct of optimal policy, relying on explicit welfare criteria Beetsma and Uhlig (1999); Alesina and Barro (2002). These criteria are typically micro-founded within New Keynesian models featuring representative agents (see, for instance, Benigno 2004; Kollmann 2004; Benigno and López-Salido 2006; Beetsma and Jensen 2005; Corsetti 2008; Gali and Monacelli 2008; Galí and Monacelli 2016; Farhi and Werning 2017; Hettig and Müller 2018; Groll and Monacelli 2020).

The present paper belongs to a new set of studies that explicitly accounts for within-country heterogeneity when revisiting open-economy issues. In particular, several studies rely on small open-economy HANK models to reassess the merits of alternative exchange-rate policies. de Ferra et al. (2021) find that household heterogeneity rationalizes "fear of floating" in the face of sudden stops. Auclert et al. (2021), in turn, stress that household heterogeneity can amplify the real income channel of exchange rates, potentially giving rise to contractionary depreciations. Guo et al. (2020) find that fixing the exchange rate leads to larger spillovers of foreign shocks but dampens their distributional impact, in contrast to what we find for HANK^2. Oskolkov (2023) and Zhou (2021) also study the distributional impact of foreign shocks and exchange-rate policies in small open-economy HANK models. Aggarwal et al. (2023) study the implications of fiscal deficits through the lens of a multi-country HANK model. Bellifemine et al. (2023) develop a HANK model of a monetary union composed of small open economies. What sets our paper apart is the two-country structure of HANK^2: it allows us to study how a monetary union alters the impact of shocks along the wealth distribution— both, vertically and horizontally across borders. Bayer et al. (2022) and Chen et al. (2023) also develop a two-country HANK models and calibrate them to the EA. They focus on fiscal frameworks rather than on monetary union as such. In Bayer et al. (2022), in particular, we develop the notion that "attitudes" towards fiscal policy may be traced back to how differences in income and wealth interact with different social security systems.

3.2 The Model

We develop a two-country New Keynesian model with incomplete markets, idiosyncratic risk, and heterogeneous agents (HANK^2). In this section, we first introduce a smaller model, a one-asset-HANK^2 model, for which we are able to establish a number of closed-form results in Section 3.3. We extend the model in Section 3.4 to a two-asset, medium-scale-HANK^2 model and calibrate it to data for the EA in order to assess the quantitative relevance of our results.

We borrow our two-country framework from Corsetti et al. (2012), while the specification of the household problem follows the small-open economy setup of Auclert et al. (2021). Their setup, in turn, extends Galí and Monacelli (2005) by allowing for household heterogeneity. Countries are isomorphic and our exposition focuses on the domestic economy or "Home". "Foreign" looks the same. Countries differ only in terms of shocks and in terms of size: We normalize the total population to unity, a fraction n of which resides in Home. In what follows, we denote foreign variables with the superscript * and use subscripts H and F to distinguish between domestic and foreign variables within a country. To benchmark the case of a monetary union against a scenario of independent monetary policies, we allow Home and Foreign to operate different currencies. In case there is a monetary union there will be an irrevocable conversion rate. We further assume that households and firms have perfect foresight and focus on a first-order approximation around the stationary equilibrium.

3.2.1 Households

There is a continuum of households, each of which faces idiosyncratic income risk. This, in turn, is due to idiosyncratic productivity, $e_{i,t}$, which is determined exogenously by a first-order Markov chain with mean $\mathbb{E}e_{i,t} = 1$. Households save via a riskless bond which is denominated in domestic currency and issued by a mutual fund which, in turn, holds government debt as well as foreign-currency bonds. This yields the familiar UIP condition but is otherwise inconsequential for the household savings decisions given perfect foresight.

Household labor supply, N_t, is determined by a labor union as described below and we assume that the labor union allocates hours worked uniformly across households. In what follows we state the household problem recursively, using time subscripts only for aggregate variables exogenous to the individual decision problem (and the value function because it is time-dependent in the face of aggregate shocks). At time t, a generic household with bond holdings a and productivity level e chooses consumption, c, and savings a', by solving the dynamic program

$$V_t(a, e) = \max_{c,a'} u(c, N_t) + \xi_t \beta E_t[V_{t+1}(a', e')] \tag{3.1}$$

$$\text{s.t. } c + a' = (1 + r_t^b)a + e\frac{W_t}{P_t}N_t - \tilde{\tau}_t e \tag{3.2}$$

$$a' \geq \underline{a},$$

P_t is the consumption price index specified below, r_t^b is the return on the bond, W_t the nominal wage, $0 < \beta < 1$ is the time discount factor, E_t the expectation operator, $\underline{a}$ an exogenous borrowing limit, and $\tilde{\tau}_t$ is a non-distortionary tax on households (it depends on e which is exogenous).[27] In addition, ξ_t is an impatience shock which we use to showcase how

[27]Non-distortionary taxes simplify our arguments below but they are not necessary: our result also holds if taxes are standard distortionary labor taxes.

country-specific demand shocks will enter the IS relation later.

For now, we assume the functional form

$$u(c, N_t) = \frac{c^{1-\gamma}}{1-\gamma} - \psi \frac{N_t^{1+\varphi}}{1+\varphi},$$

where $\gamma, \varphi > 0$ are the inverse of the intertemporal elasticity of substitution and the inverse of the Frisch elasticity of labor supply respectively.

We state the solution to the household's problem in sequence-space form as in Auclert et al. (2021), which we will use later to characterize the dynamics of the model. In particular, the solution to household's i consumption-savings problem described by (3.1) and (3.2) maps the time paths of wages, written in boldface to indicate vectors, $\boldsymbol{W}$, hours worked $\boldsymbol{N}$, real returns $\boldsymbol{r^b}$, taxes $\tilde{\boldsymbol{\tau}}$, prices $\boldsymbol{P}$, and shocks $\boldsymbol{\xi}$ to that of consumption of household i:

$$\boldsymbol{c}_i = \mathcal{C}_i\Big(\boldsymbol{W}/\boldsymbol{P}, \boldsymbol{N}, \boldsymbol{r^b}, \tilde{\boldsymbol{\tau}}, \boldsymbol{\xi}\Big). \tag{3.3}$$

Aggregating across all domestic households, we obtain an aggregate domestic consumption function $\mathcal{C}(\cdot)$, similar as in Auclert et al. (ming) or McKay and Wolf (2022):

$$\boldsymbol{c} = \mathcal{C}\Big(\boldsymbol{W}/\boldsymbol{P}, \boldsymbol{N}, \boldsymbol{r^b}, \tilde{\boldsymbol{\tau}}, \boldsymbol{\xi}\Big). \tag{3.4}$$

In each period, households allocate their consumption expenditures, c, across a domestically produced good c_H and an imported good c_F so as to enjoy overall consumption level

$$c = \left\{[1-(1-n)\alpha_H]^{\frac{1}{\sigma}} c_H^{\frac{\sigma-1}{\sigma}} + [(1-n)\alpha_H]^{\frac{1}{\sigma}} c_F^{\frac{\sigma-1}{\sigma}}\right\}^{\frac{\sigma}{1-\sigma}}. \tag{3.5}$$

Here $\alpha_H \in [0,1]$ indicates a home bias in consumption: The weight of the domestic good in total consumption is larger than what the size of the domestic economy would imply. If $\alpha_H = 1$ there is no home bias. σ is the elasticity of substitution between the domestic and the imported good. Letting $P_{H,t}$ and $P_{F,t}$ denote the price of these goods, both expressed in domestic currency, expenditure minimization implies for the consumer price index:

$$P_t = \left\{[1-(1-n)\alpha_H]P_{H,t}^{1-\sigma} + [(1-n)\alpha_H]P_{F,t}^{1-\sigma}\right\}^{\frac{1}{1-\sigma}}. \tag{3.6}$$

The optimal intratemporal allocation of expenditures implies the demand functions:

$$c_H = (1-(1-n)\alpha_H)\left(\frac{P_{H,t}}{P_t}\right)^{-\sigma} c, \quad c_F = (1-n)\alpha_H \left(\frac{P_{F,t}}{P_t}\right)^{-\sigma} c.$$

Let $\mathcal{E}_t$ denote the nominal exchange rate, that is, the price of foreign currency expressed in terms of the domestic currency. We assume that the law of one price holds, that is, the foreign currency price of the domestically produced good is given by $P^*_{H,t} = \mathcal{E}_t P_{H,t}$ and likewise for the foreign-currency price of the imported good. For future reference, it is also useful to define the terms of trade as the relative price of foreign goods to domestic goods $s_t = P_{F,t}/P_{H,t}$ and the real exchange rate $Q_t = P_t \mathcal{E}_t / P^*_t$.

3.2.2 Production

The production function is linear in labor:

$$Y_t = N_t, \tag{3.7}$$

where N_t is the aggregate labor input. For now, we assume perfect competition in the domestic goods market such that the price of domestic goods is equal to marginal costs given by the nominal wage: $P_{Ht} = W_t$. It is convenient to rewrite the real wage as a function of the terms of trade:

$$w_t = \frac{W_t}{P_t} = \frac{P_{Ht}}{P_t} = [(1-(1-n)\alpha_H) + ((1-n)\alpha_H) s_t^{1-\sigma}]^{-\frac{1}{1-\sigma}}. \tag{3.8}$$

Aggregate labor is composed of differentiated types:

$$N_t = \left(\int_k N_{k,t}^{\frac{\epsilon_t - 1}{\epsilon_t}} \right)^{\frac{\epsilon_t}{\epsilon_t - 1}}, \tag{3.9}$$

where ϵ_t is the elasticity of substitution between labor types and may vary over time. We use this "cost-push shock" to showcase how country-specific shocks to the Philips curve affect the dynamics of the model. Labor types, in turn, are efficiency units of work: $N_{k,t} = \int e_{i,t} n_{i,k,t} di$, where i indexes a household, as before, and $k \in [0,1]$ indexes the labor type. As in the recent literature, we assume that the number of hours a household works as type k, $n_{k,t}$, is determined by a union, which also determines the wage for each type $W_{k,t}$ Erceg et al. (2000); Auclert et al. (ming); McKay and Wolf (2022). A union can reset the wage with a constant probability θ.

The solution to the union problem yields a standard linearized open-economy Philips curve:

$$\hat{\pi}_{H,t} = \kappa((1-n)\alpha_H \hat{s}_t + \varphi \hat{Y}_t + \gamma \hat{C}_t) + \beta \hat{\pi}_{H,t+1} + \psi \hat{\epsilon}_t, \tag{3.10}$$

where $\pi_{H,t} := \frac{p_{H,t}}{p_{H,t-1}}$ is gross domestic producer price inflation, $\kappa \equiv \frac{(1-\theta)(1-\beta\theta)}{\theta}$, C_t is aggregate consumption, and $\psi \equiv -\frac{\kappa}{(\bar{\epsilon}-1)}$; and a "^" denote the log deviation of a variable from its steady state value.[28]

[28] We assume that the union neglects the impatience shock of households when setting wages in order to study the role of distinct demand and supply shocks.

3.2.3 Financial Markets

There are two bonds, a home bond and a foreign bond, each denoted in the country's own currency (which is identical in case there is monetary union). In the absence of arbitrage, the expected returns on both bonds are equal which implies the standard uncovered interest parity (UIP) condition:

$$1 + i_t = (1 + i_t^*) \frac{\mathcal{E}_{t+1}}{\mathcal{E}_t}, \tag{3.11}$$

with i_t being the nominal interest rate. In order to abstract from potentially heterogeneous household portfolios we assume that bond trading takes place via the mutual fund. Generally, up to the first order, the fund's portfolio is indeterminate. We assume that, in the steady state, the fund only holds domestic-currency debt. In this way, we rule out valuation effects that may arise in response to shocks under flexible exchange rates. Off steady-state, without loss of generality, we assume that cross-border trade is restricted to the foreign bond.

The domestic real interest rate is then pinned down by the Fisher equation and given by:

$$1 + r_t = \frac{1 + i_t}{1 + \pi_{t+1}}, \tag{3.12}$$

where $\pi_{t+1} := \frac{P_t}{P_{t-1}}$ is domestic CPI inflation. Given our assumption that the fund only holds domestic-currency debt in steady state, $r_t^b = r_t$ holds up to first-order.

3.2.4 Monetary and Fiscal Policy

In case there is a monetary union the common central bank adjusts interest rates based on the following simple rule:

$$i_t = \theta_\pi \Big(n\pi_{Ht} + (1-n)\pi_{Ft}^* \Big). \tag{3.13}$$

where π_{Ht} is producer price inflation at Home while π_{Ft}^* is producer price inflation in Foreign. The coefficient $\theta_\pi \geq 0$ governs the extent to which the central bank adjusts the policy rate in response to average inflation in the monetary union. In the expression above, we assume that the rate is adjusted in response to producer price inflation but in our setup "targeting" the CPI inflation is equivalent. Our results below also extend to the case where interest rates are adjusted to the (average) output gap. Note also that UIP (3.11) implies that $i_t = i_t^*$ once the nominal exchange rate is irrevocably fixed.

Alternatively, we consider a case with flexible exchange rates, assuming the following rules for monetary policy in Home:

$$i_t = \theta_\pi \pi_{Ht}, \tag{3.14}$$

and symmetrically for Foreign:

$$i_t^* = \theta_\pi \pi_{Ft}^*. \tag{3.15}$$

The conduct of fiscal policy is independent of whether there is a monetary union or not. It is set at the national level. Each government issues government bonds B_t to finance deficits and sets the tax rate. The budget constraint of the national fiscal policy reads as follows:

$$\frac{1+i_t}{\pi_t} B_t = B_{t+1} + \tau_t w_t N_t, \tag{3.16}$$

with $\tau_t = \frac{\tilde{\tau}_t}{w_t N_t}$. We assume that tax rates adjust to stabilize the level of government debt:

$$\frac{\tau_t}{\bar{\tau}} = \left(\frac{B_{t+1}}{\bar{B}}\right)^{\gamma_B^\tau}, \tag{3.17}$$

where γ_B^τ governs the speed with which debt returns to its target value $\bar{B}$.

3.2.5 Market Clearing

Bond market clearing requires:

$$A_{t+1} = B_{t+1} + \frac{B_{F,t+1}}{Q_t}, \tag{3.18}$$

that is, the total amount of domestic savings, $A_{t+1} \equiv \int_0^n a_{i,t+1} di$, equals the domestic bonds plus the net foreign asset position, $B_{F,t+1}$, which is held in foreign bonds. Analogously bond market clearing requires for Foreign:

$$A_{t+1}^* = B_{t+1}^* - \frac{n}{1-n} B_{F,t+1}. \tag{3.19}$$

Aggregating over the domestic households' budget constraints gives the net amount of domestic holdings of foreign bonds, $B_{F,t}$:

$$w_t Y_t - T_t + (1+r_t) B_t + \frac{(1+r_t^*)}{Q_t} B_{F,t} = C_t + B_{t+1} + \frac{B_{F,t+1}}{Q_t}. \tag{3.20}$$

Finally, goods markets clearing requires:

$$Y_t = \left(p_{Ht}\right)^{-\sigma} \Big[(1-(1-n)\alpha_H) C_t + (1-n)\alpha_H Q_t^{-\sigma} C_t^*\Big] \tag{3.21}$$

$$Y_t^* = \left(p_{Ft}^*\right)^{-\sigma} \Big[n\alpha_H Q_t^\sigma C_t + (1-n\alpha_H) C_t^*\Big]. \tag{3.22}$$

In Appendix C.1 we provide a formal definition of a linearized perfect-foresight equilibrium for which we derive results in the following section.

3.3 Closed-Form Results

In this section, we derive our main results in closed form. In particular, we show that a monetary union shifts the impact of country-specific shocks at the household level horizontally, that is, across borders within the brackets of the wealth distribution. We abstract from union-wide shocks because, given our assumptions on symmetry, the one-size-fits-all problem arises only in the face of country-specific shocks.

To set the stage, we first derive two propositions that show that a monetary union makes no difference for union-wide outcomes, both in terms of how aggregate variables respond to country-specific shocks and how the impact of the shock spreads vertically across the brackets of the union-wide wealth distribution. Put differently, it is irrelevant to union-wide outcomes whether countries form a monetary union or not. At the same time, a monetary union alters the adjustment to country-specific shocks across borders—both, at the aggregate level and at the household level.

For what follows, we define union-wide variables as a weighted average of the realizations in Home and Foreign, $X_t^W = nX_t + (1-n)X_t^*$, and write the canonical form for union-wide dynamics using the sequence-space representation (see Appendix C.2.1 for details).[29] As with the textbook representative agent version of the New Keynesian model, the canonical form is sufficient to describe the aggregate dynamics of the economy. Specifically, we summarize inflation dynamics with a union-wide New Keynesian Phillips curve:

$$\widehat{\boldsymbol{\pi}^{\boldsymbol{W}}} = \kappa \widehat{\boldsymbol{y}^{\boldsymbol{W}}} + \beta \widehat{\boldsymbol{\pi}^{\boldsymbol{W}}_{\boldsymbol{+1}}} + \psi \widehat{\boldsymbol{\eta}^{\boldsymbol{W}}}, \tag{3.23}$$

where η^W is a sequence of cost shocks. The union-wide IS relation, in turn, is given by:

$$\widehat{\boldsymbol{y}^{\boldsymbol{W}}} = \tilde{\mathcal{C}}_y \widehat{\boldsymbol{y}^{\boldsymbol{W}}} + \tilde{\mathcal{C}}_i \widehat{\boldsymbol{i}^{\boldsymbol{W}}} + \tilde{\mathcal{C}}_\pi \widehat{\boldsymbol{\pi}^{\boldsymbol{W}}} + \tilde{\mathcal{C}}_\xi \widehat{\boldsymbol{\xi}^{\boldsymbol{W}}}. \tag{3.24}$$

Importantly, (3.23) and (3.24), hold independently of whether there is a monetary union or not. To close the model, we need to specify a rule that pins down the union-wide interest rate $\widehat{\boldsymbol{i}^{\boldsymbol{W}}}$. This is where the monetary union comes into play. However, given that the interest-rate rule for the union has the same functional form as the rules under monetary independence—except that it targets weighted average inflation—it is irrelevant for the dynamics of the union-wide interest rate whether or not the two countries operate a monetary union.

Proposition 2 *The union-wide aggregate dynamics are characterized by* (3.23) *and* (3.24) *and a mapping from aggregate union-wide inflation to aggregate union-wide policy rates. Because under the assumptions in Section 3.2.4 above this mapping is the same in a monetary union and with independent monetary policies, so are aggregate dynamics.*

[29] For lack of a better term, we also refer to these variables as "union-wide" variables even if the two countries operate independent monetary policies and let the exchange rate float.

With a monetary union, we have $i_t = i_t^* = i_t^W$ and given equation (3.13), we have:

$$i_t^W = \theta_\pi(n\pi_{H,t} + (1-n)\pi_{F,t}^*) = \theta_\pi \pi_t^W.$$

With independent monetary policies, given by equations (3.14) and (3.15), the union-wide interest rate is:

$$i_t^W = ni_t + (1-n)i_t^* = n(\theta_\pi \pi_{H,t}) + (1-n)(\theta_\pi \pi_{F,t}^*) = \theta_\pi \pi_t^W \tag{3.25}$$

and, hence, *exactly the same as with a monetary union.*

Proposition 2 implies that if a monetary union experiences country-specific shocks, union-wide aggregates like output, consumption, and inflation behave exactly the same independently of whether countries form a monetary union or not. To see why, consider a shock originating in Foreign. In a monetary union, the response of monetary policy is a response to the weighted average of the dynamics in both countries. This implies, for instance, that monetary policy reacts "too much" from the perspective of Home and "too little" from the perspective of Foreign, compared to what would happen under independent policies. But given that the countries are isomorphic—in particular given $(\kappa, \mathcal{C}, \theta_\pi) = (\kappa^*, \mathcal{C}^*, \theta_\pi^*)$—"too little" and "too much" means the same in absolute value and, thus, the contribution of each country to union-wide dynamics exactly offsets each other. Note that this holds even if the countries are not of the same size, in which case the size-weighted absolute value would be the same.

From a union-wide perspective, monetary union is also irrelevant to the impact of shocks along the wealth distribution. To see this, consider a generic household j in Home. Given symmetry, there are $\frac{n}{1-n}$ times identical households in Foreign, that is, households with the same idiosyncratic productivity and the same wealth. We label these twin households j^* and define $c_J = nc_j + (1-n)c_{j^*}$ as aggregate consumption of household j and its twins. Note that in linearized form, we have

$$\widehat{\boldsymbol{c}}_j = \mathcal{C}_{w,j}\widehat{\boldsymbol{w}} + \mathcal{C}_{N,j}\widehat{\boldsymbol{N}} + \mathcal{C}_{i,j}\widehat{\boldsymbol{i}} + \mathcal{C}_{\pi,j}\widehat{\boldsymbol{\pi}} + \mathcal{C}_{\tau,j}\widehat{\boldsymbol{\tau}} + \mathcal{C}_{\xi,j}\widehat{\boldsymbol{\xi}} \tag{3.26}$$

$$\widehat{\boldsymbol{c}^*}_{j^*} = \mathcal{C}^*_{w^*,j^*}\widehat{\boldsymbol{w}^*} + \mathcal{C}^*_{N^*,j^*}\widehat{\boldsymbol{N}^*} + \mathcal{C}^*_{i^*,j^*}\widehat{\boldsymbol{i}^*} + \mathcal{C}^*_{\pi^*,j^*}\widehat{\boldsymbol{\pi}^*} + \mathcal{C}^*_{\tau^*,j^*}\widehat{\boldsymbol{\tau}^*} + \mathcal{C}^*_{\xi^*,j^*}\widehat{\boldsymbol{\xi}^*}. \tag{3.27}$$

Given symmetry, $\mathcal{C}_{x,j} = \mathcal{C}^*_{x^*,j^*}$ and, thus:

$$\widehat{c}_J = \mathcal{C}_{w,j}\widehat{\boldsymbol{w}^{\boldsymbol{W}}} + \mathcal{C}_{N,j}\widehat{\boldsymbol{N}^{\boldsymbol{W}}} + \mathcal{C}_{i,j}\widehat{\boldsymbol{i}^{\boldsymbol{W}}} + \mathcal{C}_{j,\pi}\widehat{\boldsymbol{\pi}^{\boldsymbol{W}}} + \mathcal{C}_{j,\tau}\widehat{\boldsymbol{\tau}^{\boldsymbol{W}}} + \mathcal{C}_{\xi,j}\widehat{\boldsymbol{\xi}^{\boldsymbol{W}}}. \tag{3.28}$$

The same logic applies to all policy functions of the household. Given Proposition 2, the inputs of the aggregate policy functions of the twin households do not depend on whether there is a monetary union or not. Hence, the weighted average (or union-wide aggregate) of a choice variable of household j in Home and its $\frac{n}{1-n}$ twins in Foreign does therefore not depend on whether there is a monetary union or not. Our next irrelevance result follows directly:

Proposition 3 *The impact of country-specific shocks along the union-wide wealth and income distribution is independent of whether two countries form a monetary union or not. In other words, the monetary union does not alter the impact of the shock vertically.*

Against this background, the next proposition follows directly. It summarizes our main result.

Proposition 4 *Monetary union shifts the distributional impact of country-specific shocks horizontally across borders within the brackets of the wealth distribution.*

To see what drives this result, recall from our arguments above that a monetary union alters the dynamics of *country-specific* variables relative to what would be observed under independent monetary policies. This means that the arguments that feature in the consumption function of individual households in Home and Foreign, (3.26) and (3.27), generally differ compared to what would be the case with independent monetary policies. Also, the consumption choice of a household with a given wealth and productivity state in Home will generally differ from that of its twin in Foreign. Yet, as established in Proposition 3, how the *union-wide* wealth distribution changes in response to country-specific shocks does not depend on the monetary union (because it does not shift the impact of the shock vertically). Assuming countries are of the same size, this then requires that monetary union changes the effect of a shock on a household's consumption choice in Home in exactly the opposite way as it does for its Foreign twin. When countries differ in size, the differential impact of monetary union on the choice of a generic household in Home is of the opposite sign as that of its Foreign twin, weighted by the number of twins that a Home household has in Foreign. It follows that the distributional effect of monetary union operates horizontally across borders: It shifts the distributional impact of shocks (compared to a scenario of independent monetary policy) between households in Home and Foreign *with the same individual states* or within the same bracket of the wealth distribution. For instance, if consumption of the poor at Home is higher with a monetary union in place than with independent monetary policies after a given shock, consumption of the poor in Foreign must be lower by the (weighted) same amount.

3.4 Quantitative Analysis

Our analysis has established that a monetary union alters the impact of business cycle shocks at the household level. It does so by shifting the adjustment horizontally across countries within the brackets of the wealth distribution. We now perform a quantitative analysis in order to assess how strongly this effect plays out for different types of households. We perform the quantitative analysis in a version of the model that is extended along a number of dimensions and, importantly, it no longer restricts Home and Foreign to be symmetric at the household level. Specifically, we calibrate the model to two countries of the euro area that represent polar cases in terms of the wealth distribution: Germany and Italy. For this version of the model, we also show that the results established by Propositions 2 and 3, which rely on symmetry, are still approximately satisfied.

3.4.1 Medium-sized HANK^2 model

Since our question at hand is a quantitative one, we enrich our model laid out in Section 3.2 by features that are frequently used in medium-sized business cycle models. In particular, we use a two-country version of the model developed in Bayer et al. (ming) which has been shown to be able to generate business-cycle dynamics that conform well with the data. We calibrate this medium-sized HANK^2 model to capture key aspects when it comes to asset holdings and wealth distributions in Germany and Italy. In what follows, we briefly sketch the main extensions of the model and delegate a full description of the extended model to Appendix C.3. As before, the structural features are the same in Home and Foreign. Yet by assigning different parameter values below we make sure that Home and Foreign differ—in accordance with the data.

Households. We modify the household side in three ways in order to be better able to match the wealth distribution in the data. First, we assume that a group of households is employed by firms while others are self-employed entrepreneurs. The former group receives only labor income while entrepreneurs earn firm profits that arise due to monopolistic competition in the goods market (see below). Yet, households may move from one group (or employment state) to the other according to some exogenous probability. Both labor income and profit income are subject to a proportional income tax. Second, we assume that households can hold two different types of assets, liquid government bonds, and illiquid capital. Capital holdings are illiquid because we assume that only a random share of households can trade capital in a given period. Third, we assume that in Foreign, which will be calibrated to Germany, households will receive a minimum income benefit which we model as a targeted transfer which those households receive whose income is below a certain threshold. As Bayer et al. (2022) show in detail, large differences in minimum income benefits across Germany and Italy can explain a large part of the differences in the wealth distribution,[30] and, as a result, requires large differences in government debt (high in Italy and low in Germany) in order to obtain the same real interest in both countries in steady state.

Firm sector. We also extend the firm sector by assuming that not only wages but also prices are adjusted infrequently. To this end, we assume a multi-layered production structure.

[30] Pham-Dào (2016) shows that this is the case more generally across euro area countries. However, she uses a single-asset incomplete markets model.

Intermediate goods producers operate under perfect competition using both domestic capital and labor which we assume are immobile across countries. We also assume that production is subject to country-specific total factor productivity (TFP). Final good producers, in turn, differentiate domestic intermediate goods under monopolistic competition and are subject to Calvo (1983)-type price setting frictions in Home and Foreign. Domestically and imported goods are then bundled into consumer goods as in Section 3.2. Capital producers also use intermediate goods and face quadratic investment adjustment costs.

Fiscal policy. Lastly, we also consider a somewhat richer set of fiscal policies. First, the government in Foreign has to fund the minimum income benefits. Second, we now also consider government spending. This will allow us to analyze how a government spending shock plays out, both under flexible exchange rates and in the monetary union.

Shocks. In what follows we focus on TFP shocks and government spending shocks which may originate either in Home or Foreign. We assume each of these four shocks follows an exogenous AR(1)-process.

3.4.2 Symmetric Calibration

Compared to the stylized model in Section 3.2, the medium-sized HANK^2 model outlined here features a richer structure in order to better capture key aspects of the data, both at the micro and the macro level. Before we calibrate the model to the EA, we therefore verify that the results established by Propositions 2 - 4 still hold exactly once we simulate a perfectly *symmetric* version of the model.

For this purpose, we pick parameter values for both Home and Foreign in line with the "Italy calibration" below. We provide detailed results in Appendix C.5. We study, in particular, the transmission of TFP shocks at the country level and at the union level. We find that monetary union alters the effects of a country-specific shock at the country level very much. Yet, while there is "too little" adjustment in one country and "too much" in the other these effects also offset each other completely in the larger model—in line with Proposition 2. The response of union-wide prices and quantities to country-specific shocks is independent of whether the countries operate a monetary union or not. Consequently, as argued above, the terms that enter the consumption function of union-wide twins are the same (Proposition 3), implying that monetary union does not shift the impact of shocks vertically across the wealth and income distribution. Instead, it shifts the impact of shocks horizontally across countries, as established in Proposition 4. To the extent that our results below differ from those stated in Propositions 2 - 4, this thus reflects the asymmetric calibration which in turn captures the differences in household-level heterogeneity in Germany and Italy.

3.4.3 Asymmetric Calibration to the Euro Area

We outline how we calibrate the model to the EA and refer readers to Appendix C.4 for more details. Importantly, we now allow countries to differ not only in terms of shocks but also in terms of heterogeneity at the household level, in line with the data for Germany and Italy. For this purpose, we set parameters to target the wealth distributions and asset holdings in both countries.

For most parameters, we use standard values as listed in Appendix C.4 We specify the parameters that determine the income process at the household level to match micro-level estimates established for German and Italian data. In particular, we set the persistence of idiosyncratic income shocks to a standard value found for the euro area, see for example Pham-Dào (2016); and set the respective standard deviations as to match income inequality in Italy and in Germany. Moreover, we assume that in Foreign there are minimum income benefits which, overall, amount to 1% of GDP, in line with data for Germany. There are no minimum income benefits in Home which represents Italy.[31]

We then use six parameters to target key features of the wealth distributions and asset holdings in Germany and Italy. In particular, we use the discount factor, the portfolio adjustment probability, the probability which governs the transition of households to become entrepreneurs, and the borrowing penalty to match the level of government debt, the capital-to-output ratio, the wealth Gini, the top-10% wealth share, the bottom-50% wealth share, and the mass of borrowers.

Table A.3 reports key moments as predicted by the model under the baseline calibration and contrasts them with their empirical counterparts. The top panels show values for the steady state where the empirical moments have been used as calibration targets. Note that the model is able to generate the observed large asymmetry between both countries: The wealth distribution is much more unequal in Germany compared to Italy, while government debt is considerably higher in Italy. As explained in detail in Bayer et al. (2022), from the perspective of the model these two aspects are interrelated and can be explained by the stronger need for self-insurance when minimum income benefits are not available. It should be noted that the degree of asymmetry between the two countries in these areas is one of the largest in the euro area (see also Pham-Dào 2016; Kindermann and Kohls 2017).

At the same time, the model is able to capture key features of the business cycle. For this purpose, we set parameters which capture macro frictions and policies in line with estimates of Bayer et al. (ming). Note that for the baseline, we assume a common monetary policy (monetary union) and assume that the exchange rate is permanently fixed. Monetary policy is described by an interest rate feedback rule with interest rate smoothing. For the open economy parameters, we rely on standard parameter values in the literature, see again Appendix C.4. We follow Enders et al. (2013) in specifying country-specific government

[31]The discretization of the income process effectively also establishes a non-zero lower income bound for the Home country, that is however much lower than the Foreign country's minimum income benefits.

Table 3.1: Calibrated Model v Data

			Model		Data	
			H	**F**	**ITA**	**GER**
Steady state	Assets	Debt (% of output)	132	71	132	71
		Capital-Output-Ratio	3.3	3.2	3.3	3.2
(targeted)	Distribution	Wealth gini	0.60	0.72	0.61	0.73
		Top-10% wealth share	0.43	0.55	0.44	0.52
		Bottom-50% wealth share	0.10	0.01	0.09	0.02
		Borrowers	0.08	0.18	0.08	0.18
Business Cycle	Volatility	Std(Y)*100 (targeted)	3.78	2.90	3.78	2.74
		Std(C)/Std(Y)	0.95	0.86	0.95	0.90
		Std(I)/Std(Y)	2.52	3.00	1.82	1.60
		Std(π)/Std(Y)	0.63	0.67	0.33	0.40
	Co-Movement	Corr(Y, Y^*) (targeted)	0.80		0.80	
		Corr(C, C^*)	0.95		0.79	
		Corr(I, I^*)	0.89		0.33	
		Corr(π, π^*)	0.97		0.77	

Notes: Model predictions based on baseline calibration, see Appendix C.4 for details. Micro data based on the 2017 wave of the Household Finance and Consumption survey of the ECB. Macro data from Eurostat and Worldbank (Inflation). Quantities are measured in real per capita terms, yoy changes; sample: 1999Q1-2022Q2.

spending shocks and TFP shocks. We also add a common TFP shock since, otherwise, the model cannot account for the high degree of business-cycle comovement across both countries. We specify TFP shocks and the common component as we target output volatility in both countries and the co-movement of output across countries. The lower part of Table A.3 shows that our model then does a fairly good job in matching other key statistics of the Euro area business cycle. In particular, the relative volatility of investment, consumption, and inflation is in the right ballpark, as is the co-movement across countries.

3.4.4 Macroeconomic Adjustment to Country-specific Shocks

We use the calibrated model to analyze the macroeconomic adjustments to country-specific shocks. In particular, we consider a TFP shock originating in Foreign and a government spending shock originating in Home. Even though the countries are no longer symmetric, a very similar pattern emerges for shocks that originate in the other country. Hence we do not report results for this case to economize on space.

Figure 3.1 shows the responses of country aggregates in Home and Foreign as well as the union-wide aggregates to a contractionary TFP shock that originates in Foreign. Throughout, we contrast results for the monetary union (blue solid line) with those for independent monetary policies (red dashed line) in Home (left column), in Foreign (middle

Figure 3.1: Adjustment to adverse TFP shock originating in Foreign

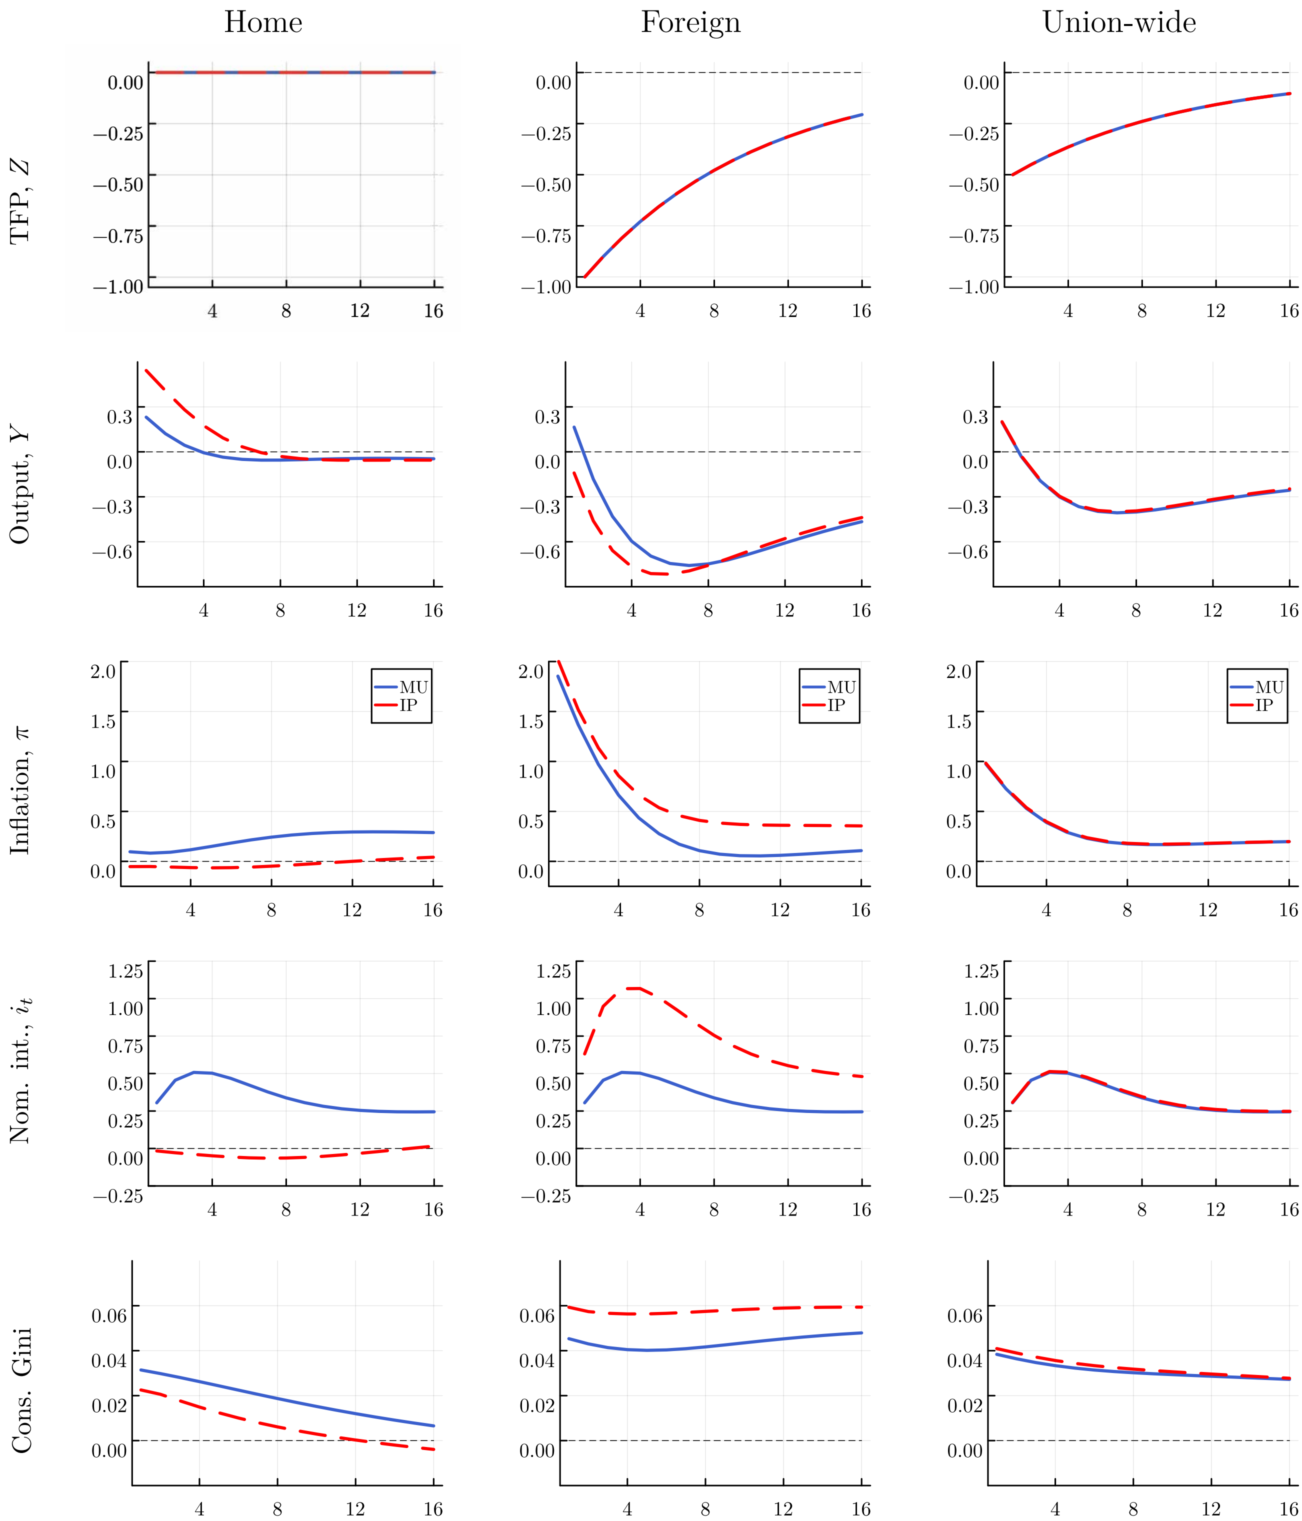

Notes: monetary union v independent monetary policies in Home (left), in Foreign (middle), and aggregate of Home and Foreign (right). Y-axis: Percentage deviation from steady state and percentage points in case of interest rates. X-axis: Quarters.

column), and in the entire union by displaying the aggregate responses (right column). Recall that in case of independent monetary policies, the interest rate feedback rule is the same as in the monetary-union case, except that monetary policy in each country responds to country-level rather than union-wide inflation rates. In each panel, the horizontal axis measures time in quarters, the vertical axis measures the percentage (or percentage point) deviation from steady state.

The top panels show the shock process which is independent of whether there is a monetary union or not: TFP in Foreign contracts; it is unchanged in Home. The second row shows the adjustment of output. Here monetary union makes a fundamental difference. Output in Foreign *increases* on impact with a monetary union in place, but *decreases* under independent policies. Likewise, output in Home also responds very differently across monetary regimes: it increases much more under independent monetary policies.

To rationalize these differences, it is instructive to study the adjustment of the policy rate, shown in the fourth row: in the monetary union it responds in the same way in both countries, while with independent monetary policies, we observe an increase in the short rate in Foreign and a decline in Home. This reflects, in turn, the differential impact of the shock on inflation in Foreign and Home which is shown in the third row: the contractionary TFP shock is strongly inflationary in Foreign, thus necessitating a monetary contraction. The policy response is much weaker in Foreign once it operates in a monetary union. For Home, it is the opposite: operating inside the monetary union implies a restrictive instead of an expansionary monetary policy. This is the one-size-doesn't-fit-all issue that is at the heart of the policy discussion in monetary unions. In our context, it even induces a change in the sign of the response of Foreign output in response to a TFP shock.

The right column of Figure 3.1 shows the aggregate response of Home and Foreign under monetary union and with independent monetary policies. For all variables, the two impulse responses lie almost perfectly on top of each other. This shows that even though the two countries now differ substantially in terms of household-level heterogeneity, the result of Proposition 2 still holds approximately: monetary union does change the adjustments to country-specific shocks at the country level but it does so by shifting the adjustment between countries. The overall effect of a monetary union on the adjustment of union-wide aggregates turns out to be negligible.

Against this background, the bottom panels of the figure show the response of the consumption Gini. Consumption inequality increases after the shock in both countries and across both exchange rate regimes for reasons which become clear below. At this point it is important to point out that the response of aggregate, union-wide consumption inequality (shown in the right column) is basically independent of the exchange rate regime—consistent with Proposition 3. Yet consumption inequality increases less at Foreign and more at Home in case of monetary-union, compared to what we observe under independent monetary policies.

Turning to the effects of a government spending shock, shown in Appendix C.6, we observe that monetary union also alters the dynamics within countries profoundly. In particular, a monetary union amplifies the output response in Foreign, but dampens it in Home. This again reflects the common monetary stance in the union. However, the overall effect of monetary union on the adjustment of union-wide aggregates is close to zero—just like in the case of TFP shocks. Hence, the result of Proposition 2 also holds approximately for government spending shocks, even if we allow countries to have asymmetries.

3.4.5 Adjustment at the Household Level

We are finally in a position to address the main question of the paper: How does monetary union alter the impact of country-specific shocks at the household level? Proposition 4 above establishes for the symmetric case that monetary union indeed shifts the adjustment to country-specific shocks across borders. Specifically, monetary union shifts the adjustment at the household level horizontally across borders within the brackets of the wealth distribution. For the calibrated version of the model, we have shown that the results established in Propositions 2 and 3 approximately hold in the larger model even as it is calibrated to capture cross-country heterogeneity at the household level.

We now use this version of the model to quantify how monetary union alters the impact of shocks at the household level. Specifically, we compute the welfare impact of a shock for each household across the wealth distribution and contrast results for the monetary union with those for independent monetary policies. We measure the welfare impact using the consumption equivalent variation, which is the permanent consumption change that would make an individual household equally well off as the shock under consideration. We stress upfront that we take an ex-post perspective, evaluating welfare based on specific shocks (that is, one-sided welfare), rather than providing an ex-ante welfare analysis based on a second-order approximation of the utility function.

To synthesize results, we compute how the consumption equivalent variation due the shock under consideration changes for each decile of the wealth distribution as countries move from independent monetary policy to monetary union. Figure 3.2 shows the results. The left (right) panels show the welfare differences in the adjustment due to monetary union for an adverse TFP shock that originates in Foreign (Home). The upper (lower) panel depicts the effect along the deciles of the wealth distribution in Home (Foreign). The emerging pattern is clear-cut and warrants three observations. First, the pattern is consistent with the result of Proposition 4 according to which monetary union shifts the impact of shocks across borders within the brackets of the wealth distribution. For the calibrated model, this does not hold exactly because we relax the assumption of country symmetry. And yet, we find that the adjustment in Home and Foreign is still fairly symmetric across the wealth distribution—the adjustment in Foreign mirrors those in Home. Interestingly, we observe that there is now some vertical shift in the impact of the shock: the overall effect

Figure 3.2: How monetary union alters the welfare impact of shocks

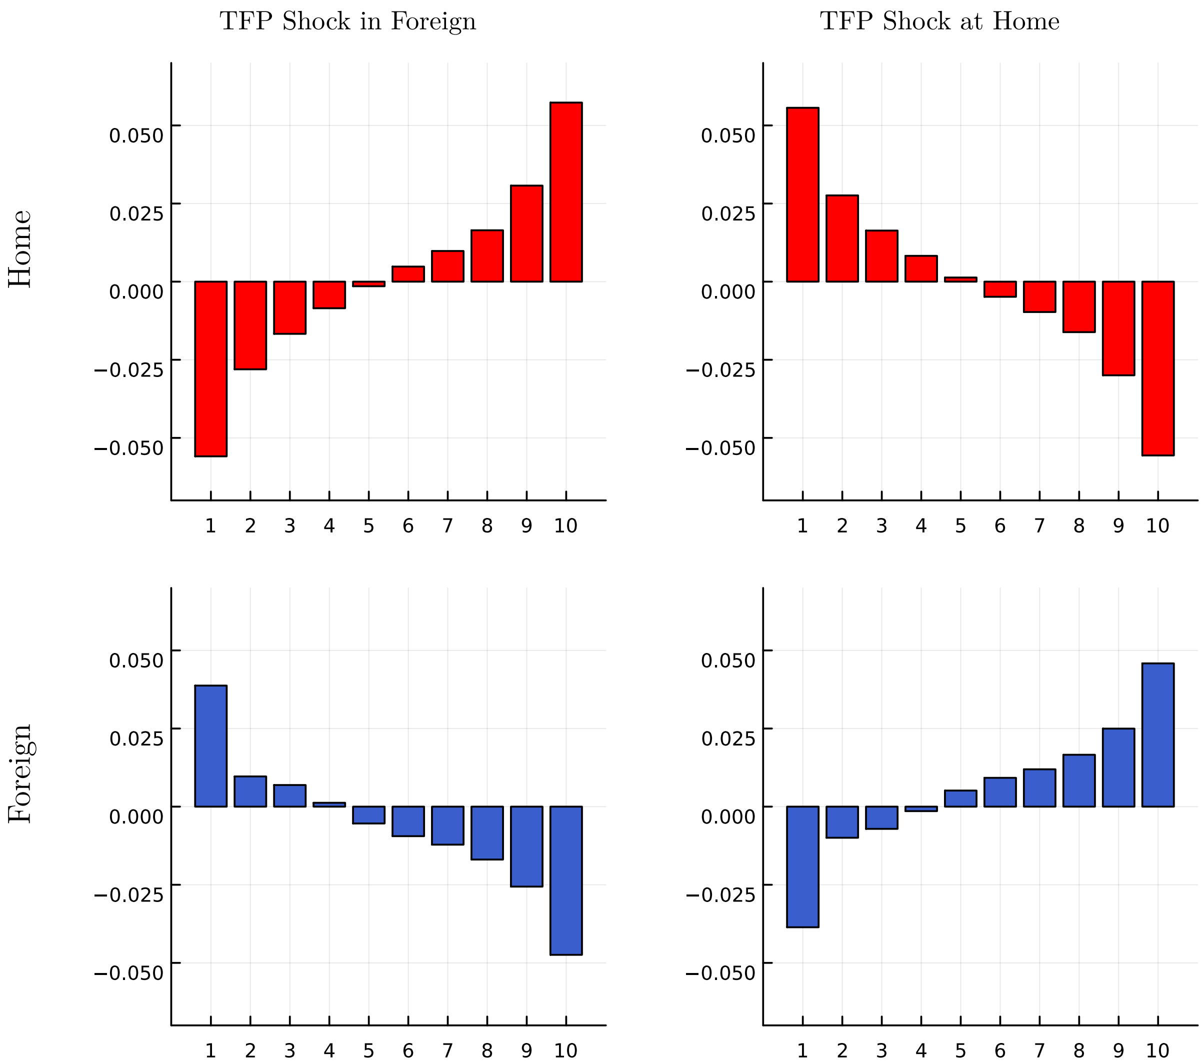

Notes: Difference of welfare impact of a Foreign adverse TFP shock (left) and a Home adverse TFP shock (right) between monetary union and independent monetary policies in Home (upper panel) and Foreign (lower panel). Y-axis: Difference in terms of consumption equivalent compensating variations. X-axis: Wealth deciles.

for each wealth bracket in Home and Foreign is not exactly zero.[32]

Second, we find that the change in the impact of shocks on welfare due to monetary union is much more concentrated in the tails of the wealth distribution. The middle class is much less affected. Put differently, whether there is a monetary union in place or not matters for the shock's welfare impact, but only for the poor and the rich. This is consistent with Proposition 4, and holds true for both Home and Foreign.

Third, these patterns do not depend on the origin of the shock. This becomes clear when comparing the left and the right column of the figure. If the shock originates in Home

[32]Figure C.1 in the Appendix reports results for the symmetric calibration: in this case, the patterns in both countries are also perfectly symmetric.

instead of Foreign, the exact same patterns emerge—just with a flipped sign. Likewise, the pattern emerges for positive and negative shocks equally; and it is not specific to TFP shocks. A similar pattern also emerges for government spending shocks (see Figure C.4 in the Appendix).

To see how monetary union alters the welfare impact of shocks along the wealth distribution, we decompose the overall effect for each decile into partial equilibrium effects. For this purpose, we exploit the fact that how the welfare of a given household is affected by a TFP shock depends on the arguments that enter its policy functions and thus the choice of variables which directly impact its welfare. By changing the adjustment of these at the country level, a monetary union changes the effect of a shock on households' welfare.

To economize on space, we focus on an adverse TFP shock originating in Foreign and show results in Figure 3.3. Based on our decomposition, each row shows the contribution of a specific variable to the overall effect: the real rate on liquid bonds, capital income, labor and profit income, and taxes. The blue bars represent the consumption equivalent variation of the shock under a monetary union, while the red bars represent the counterpart for the case of independent monetary policies. The left (right) panels report results for the deciles of the wealth distribution in Home (Foreign). We observe that monetary union has a strong bearing on the welfare impact of the shock by changing the way the real interest rate (top panel) and taxes (bottom panel) respond to the shock, and more so than for labor and capital income (middle panels). This is intuitive because—as discussed above—monetary union changes the interest rate response to a country-specific TFP shock. Changes in the interest rate then impact governments' budgets by altering the interest rate burden on the outstanding debt which, ultimately, results in an adjustment of the tax rate.

Importantly, the way in which the different adjustments of the real interest rate and the taxes affect households' welfare is highly heterogeneous along the wealth distribution. And we observe that the changes due to monetary union are largest at the tails of the wealth distribution—reflecting a different interest rate exposure. High-wealth households are directly exposed to interest rate changes through their assets, while low-wealth households are exposed through the tax response. As shown in Figure 3.1 above, a negative TPF shock in Foreign raises interest rates and hence the return on the liquid asset in Home. At the same time, wages fall and labor income taxes rise. This benefits the asset-rich and harms the asset-poor. The monetary union changes the size of these price responses. As a result, the welfare impact of the shock on the poor and the rich depends on whether there is a monetary union in place or not.

Figure 3.3 also illustrates why the middle class is largely unaffected by a monetary union. In this case monetary union also changes the welfare impact of different aspects of the shock: for example, the middle class in Home benefits from the higher real rates in response to an adverse TFP shock in Foreign and it suffers from the increased tax response. But these two effects roughly cancel each other out. This is intuitive as the middle class roughly holds an average amount of wealth and pays taxes on roughly an average income.

Figure 3.3: Decomposition of welfare effect of TFP shock in Foreign

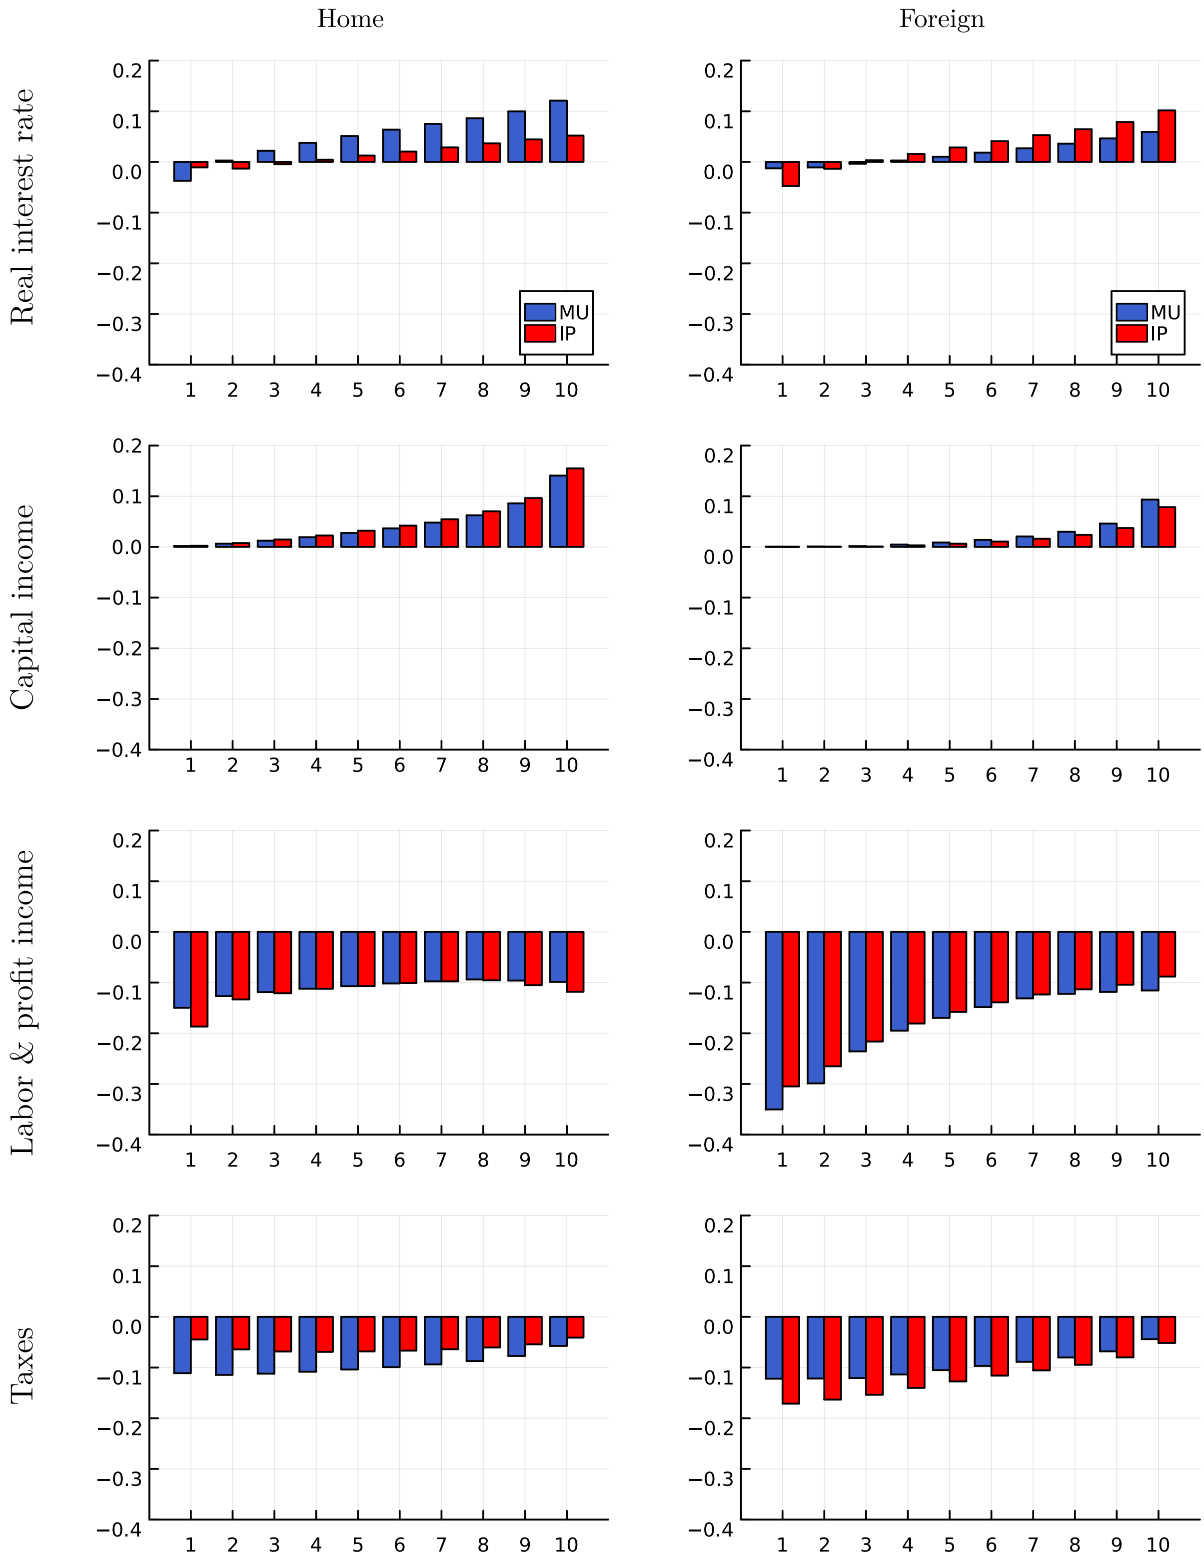

Notes: Decomposition of welfare effects of a contractionary TFP shock in Foreign across the deciles of the wealth distribution. Monetary union (blue) vs Independent monetary policies (red). Y-axis: Difference in terms of consumption equivalent compensating variations. X-axis: Wealth deciles.

These results offer a new perspective on some of the policy debates surrounding the euro. In fact, they may explain why the European Monetary Union did not break up in the face of sizeable asymmetric shocks during its 20-year-plus history. In every country, those that benefit from the (union-wide) monetary response can always form a sufficiently large coalition with the middle class to support the union, as long as there is a small (and here non-modeled) cost of breaking up the union. Focusing on how monetary union alters the welfare impact of business cycle shocks thus offers new insights into the political economy of monetary unions, an issue that calls for further research.

3.5 Conclusion

Asymmetric shocks are a classic theme of OCA theory. They bring to the fore the one-size-doesn't-fit-all problem from which monetary unions are bound to suffer at times. We revisit the issue through the lens of a Heterogeneous Agent New Keynesian model with two countries: HANK^2. It belongs to a class of models that breaks with the representative agent paradigm and offers new perspectives. In particular, in contrast to earlier generations of OCA theory, we are no longer confined to analyzing what membership in a monetary union means for countries or regions as such.

Instead, we can investigate how monetary union alters the impact of shocks for individual households. In particular, the HANK^2 structure allows us to distinguish how monetary union alters the impact of shocks horizontally across borders within the brackets of the wealth distribution and vertically across the wealth union-wide brackets. A key result of our analysis is that a monetary union shifts the adjustment to shocks horizontally and not so much vertically. We show in closed form that a monetary union neither changes the union-wide dynamics after a country-specific shock nor its vertical impact across the brackets of the union-wide wealth distribution. Instead, it shifts the impact horizontally across borders within the brackets of the wealth distribution, from the poor in one country, for instance, to the poor in the other country.

Our quantitative analysis shows that this effect is particularly strong for the tails of the wealth distribution and weaker for the middle class. This brings to the fore questions about the political economy of currency unions which we take up in a companion paper Bayer et al. (2022). Here we just note that our results may provide a rationale for why the EA did not break up during its 20-year-plus history despite several severe crises and calls for an exit of individual countries: In the face of a specific shock (or crisis), whether a country operates inside a monetary union or not does not matter so much for a large fraction of the population. That being said, we find that it can matter a lot for a smaller fraction of the population. This part of the population, however, lacks political majorities.

Chapter 4

A Behavioral Heterogeneous Agent New Keynesian Model

with Oliver Pfäuti

Abstract

We analyze how cognitive discounting and household heterogeneity affect the transmission of monetary policy. Under cognitive discounting, households' expectations exhibit an under-reaction to news about the aggregate economy, which is consistent with empirical evidence on household expectations. Our model simultaneously accounts for recent empirical findings of the transmission of monetary policy: (i) monetary policy affects consumption largely through indirect effects, (ii) households are unequally exposed to aggregate fluctuations and income risk is countercyclical, (iii) forward guidance is less powerful than contemporaneous monetary policy, (iv) and the economy remains stable at the zero lower bound. In contrast to demand shocks, supply shocks are amplified through both, cognitive discounting and household heterogeneity, such that inflation increases more than twice as strong as when abstracting from cognitive discounting and household heterogeneity.

Keywords: Monetary Policy; Heterogeneous Households; Behavioral Macroeconomics; Forward Guidance; Lower Bound; Inflation; Macroeconomic Stabilization
JEL-Classification: E21, E52, E62, E71.

4.1 Introduction

Recent empirical evidence has led to a rethinking of how monetary policy is transmitted to the economy: (i) monetary policy affects household consumption to a large extent through changing people's incomes rather than directly through changes in the real interest rate. These *indirect effects* tend to amplify the effects of conventional monetary policy on consumption as (ii) the incomes of households that exhibit higher marginal propensities to consume are found to be more exposed to aggregate income fluctuations induced by monetary policy; (iii) announcements of future monetary policy changes, in contrast, have relatively weak effects on current economic activity; and (iv) advanced economies have not experienced large instabilities in times in which the nominal interest rate has been stuck at the lower bound.[33]

In this paper, we propose a new framework that accounts for these four facts *simultaneously*: the behavioral Heterogeneous Agent New Keynesian model—or *behavioral HANK model*, for short. The model features a standard New Keynesian core with nominal rigidities, but we allow for household heterogeneity and bounded rationality in the form of cognitive discounting. The presence of both—household heterogeneity and bounded rationality—is key to account for the four facts jointly. In contrast to existing models, our model accounts for the four facts without having to rely on a specific monetary or fiscal policy.

We first illustrate how cognitive discounting interacts with household heterogeneity under a specific calibration of our model for which we obtain a closed-form solution but that still captures the key features of the model. Households that exhibit higher marginal propensities to consume are more exposed to monetary policy which is crucial to account for the fact that monetary policy is amplified through indirect general equilibrium effects. Under cognitive discounting, households' expectations underreact to aggregate news—consistent with what we document for household survey expectations—which dampens the effects of announced future monetary policy changes and ensures that the model remains stable at the effective lower bound. Second, we then show numerically that all our results carry over to the full model. This holds true, even when households over- or underreact to idiosyncratic shocks or when households are heterogenous in their behavioral biases.

Accounting for these four facts simultaneously has important implications for macroeconomic stabilization. In particular, we uncover a new amplification channel of adverse supply shocks: the unequal exposure of households, their behavioral bias and the interaction of the two lead to a substantial increase in the output gap and inflation. Inflation increases more than twice as strong as when abstracting from these model features. As a consequence of this amplification channel, there is a strong trade-off for monetary policy between price

[33] See, e.g., Ampudia et al. (2018), Slacalek et al. (2020) and Holm et al. (2021) for the empirical relevance of indirect channels in the transmission of monetary policy, Auclert (2019), Patterson (2023) and Slacalek et al. (2020) for evidence on households' income exposure and their marginal propensities to consume, and see, for example, Del Negro et al. (2015), D'Acunto et al. (2022), and Roth et al. (2021) for empirical evidence on the (in-)effectiveness of monetary policy announcements about its future actions, and Debortoli et al. (2020) and Cochrane (2018) on the stability at the lower bound.

stability on the one side and fiscal and distributional consequences on the other side after an inflationary supply shock. If monetary policy wants to fully stabilize inflation, it needs to increase interest rates much more aggressively, which pushes up the government debt level and inequality more strongly.

Our model builds on the recent heterogeneous-agent New Keynesian literature (HANK) which combines the typical Bewley-Huggett-Aiyagari incomplete markets setup with nominal rigidities. Ex-ante identical households face uninsurable idiosyncratic productivity risk, incomplete markets and borrowing constraints. In contrast to that literature, households in our model do not necessarily hold rational expectations. In particular, we allow for *cognitive discounting* of aggregate variables: households anchor their expectations about future macroeconomic variables to the steady state and cognitively discount expected future deviations as in Gabaix (2020). As a result, expectations then underreact to aggregate news, as we show to be the case empirically across all income groups and which is also consistent with findings in D'Acunto et al. (2022) or Roth et al. (2021).[34]

We start by showing that for a specific calibration, the model simplifies such that it can be solved in closed form. In particular, the household block can be represented as if there were two representative households. Yet, the model still shares the key features with our full model, namely unequal exposure of households to aggregate shocks, a precautionary savings motive of households and borrowing constraints as well as cognitive discounting of aggregate shocks.[35] The two *as-if* representative households differ in the following respects: the first group is "unconstrained", in the sense that they participate in financial markets and are on their Euler equation. The second group consists of "hand-to-mouth" households who consume all their disposable income. They exhibit high marginal propensities to consume (MPCs) and their income is more exposed to monetary policy in line with the data. As unconstrained households face a risk of becoming hand-to-mouth, they exhibit a precautionary-savings motive.

Given this specific calibration, the model can then be represented in just three equations exactly like the textbook Representative Agent New Keynesian (RANK) model. The key novelty is a new aggregate IS equation. In contrast to the textbook model, our IS equation features a lower sensitivity of current output to changes in expected future output due to households' cognitive discounting and a stronger sensitivity of current output to changes in the real interest rate as households with higher MPCs are more exposed to monetary policy.

As a result of the lower sensitivity of current output to future expected output, announced policies that increase future output, such as announced future interest rate cuts, are less

[34]Angeletos and Lian (2023) show how other forms of bounded rationality or lack of common knowledge can be observationally equivalent. For further evidence on underreaction of expectations or general patterns of inattention, see, e.g., Coibion and Gorodnichenko (2015), Coibion et al. (2022) or Angeletos et al. (2021). Kučinskas and Peters (2022) and Born et al. (2022) show that even when agents overreact to micro news, they underreact to macro news.

[35]Models with a similar household structure are often referred to as TANK (Two Agent New Keynesian) models with type switching or as THANK (Tractable HANK) models (Bilbiie (2021)). We therefore refer to this special calibration of our model as *tractable* behavioral HANK model.

effective in stimulating current output. After such an announced future interest rate cut, unconstrained households want to consume more already today as they want to smooth their consumption intertemporally. Additionally, their precautionary savings motive decreases as they would be better off in case they become hand-to-mouth in the future because hand-to-mouth households benefit more from the future boom. Cognitive discounting weakens *both* of these channels and thus, explains the lower sensitivity of current output to future expected output. The farther away in the future the announced interest rate cut takes place, the smaller its effect on today's output. Hence, the model does not suffer from the *forward guidance puzzle,* which describes the paradoxical finding in many models that announced future interest-rate changes are at least as effective in stimulating current output than contemporaneous interest-rate changes (Del Negro et al. (2015), McKay et al. (2016)). In addition, our model remains determinate under an interest-rate peg and remains stable at the effective lower bound (ELB).

The second deviation from the textbook IS equation—the stronger sensitivity of current output to changes in the real interest rate—arises because households with higher MPCs are more exposed to monetary policy. An expansionary monetary policy shock increases the income of the hand-to-mouth households more than one-for-one. As these households consume all their disposable income, this leads to a stronger response of aggregate consumption than if all households would be exposed equally to monetary policy. Thus, the model features amplification of conventional monetary policy shocks due to indirect general equilibrium effects. A decomposition into direct and indirect effects shows that indeed the major share of the monetary policy transmission works through indirect effects.

We then relax our specific calibration and show that none of our results depend on it. In particular, we build on a calibration that is standard in the HANK literature extended by cognitive discounting and the unequal exposure of households to monetary policy shocks found in the data. Consequently, the model now features a non-degenerate wealth distribution and can only be solved numerically. We show that the model still accounts for facts (i)-(iv) simultaneously.

That our model simultaneously generates amplification of conventional monetary policy through indirect effects and rules out the forward-guidance puzzle is in stark contrast to rational models. Rational HANK models that generate amplification through indirect effects exacerbate the forward-guidance puzzle. Rational models that resolve the forward-guidance puzzle, on the other hand, cannot simultaneously generate amplification of monetary policy through indirect effects (see Werning (2015), Acharya and Dogra (2020), and Bilbiie (2021)).

We extend our model in several ways. First, we consider an extension in which households are heterogeneous with respect to their cognitive discounting. We find in the data that the degree of rationality is slightly positively correlated with the income of households. Introducing this into our model, we find that this extension has only minor quantitative impacts on our results, while the model continues to account for facts (i) - (iv) simultaneously. This even applies to a version in which a subgroup of households is fully rational. Second, we allow for bounded rationality also with respect to households' idiosyncratic risk. Recent empirical findings by Kučinskas and Peters (2022) and Born et al. (2022) show that even though agents' expectations underreact to aggregate shocks, they tend to overreact to idiosyncratic shocks. We show that overreaction with respect to idiosyncratic news has only a small impact on our results: the extended model also accounts for facts (i) - (iv) simultaneously and even quantitatively, the results are barely affected by introducing bounded rationality also with respect to idiosyncratic risk.

We then show that accounting for facts (i) - (iv) simultaneously, matters greatly for the model's policy implications. Many advanced economies have recently experienced a dramatic surge in inflation which is partly attributed to disruptions in production (see di Giovanni et al. (2022)). We analyze these supply disruptions by considering a negative productivity shock.[36] We uncover a novel amplification channel of these supply shocks as both—the underlying heterogeneity and bounded rationality—amplify the inflationary pressure from the supply shock and the two mutually reinforce each other: the positive output gap redistributes towards households with higher MPCs increasing the output gap further and, thus, calls for higher interest rates in each period. As households cognitively discount these higher (future) interest rates, this further increases the output gap amplifying the redistribution to high MPCs households and therefore the increase in the output gap until the economy ends up in an equilibrium with a higher output gap and higher inflation. As a consequence, inflation increases by more than twice as much as in a model without household heterogeneity and bounded rationality.

That both—the unequal exposure of households and cognitive discounting—amplify supply shocks is in stark contrast to demand shocks. In response to persistent demand shocks, the unequal exposure of households amplifies the shock whereas cognitive discounting dampens it. Consequently, our model predicts inflation and the output gap to be less responsive to persistent demand shocks but more responsive to supply shocks compared to the rational model.

The amplification channel also implies a more pronounced trade-off for monetary policy between price stability on the one side and fiscal and distributional consequences on the other side after an inflationary supply shock. If monetary policy wants to fully stabilize inflation, it needs to hike interest rates much more aggressively to counteract the amplification forces. These stronger interest-rate hikes create side effects. In particular, they have strong fiscal implications as they increase the cost of government debt, which leads to a larger increase in government debt. Furthermore, consumption inequality increases strongly. The reason is that wealthy households benefit more from higher interest rates than asset-poor households.

Related literature. The literature treats the facts (i)-(iv) mostly independent from each other. The heterogeneous-household literature has highlighted the transmission of monetary

[36] We also consider cost-push shocks as an alternative explanation for high inflationary pressure and find similar implications for monetary and fiscal policy.

policy through indirect, general equilibrium effects (Kaplan et al. (2018), Auclert (2019), Auclert et al. (2020), Bilbiie (2020), Luetticke (2021)), and proposed potential resolutions of the forward guidance puzzle (McKay et al. (2016), McKay et al. (2017), Hagedorn et al. (2019), Acharya and Dogra (2020), McKay and Wieland (2022)). Werning (2015) and Bilbiie (2021) combine the themes of policy amplification and forward guidance puzzle in HANK and establish a trade-off inherent in models with household heterogeneity: if HANK models amplify contemporaneous monetary (and fiscal) policy through redistribution towards high MPC households, they dampen precautionary savings desires after a forward guidance shock which aggravates the forward guidance puzzle.

Few resolutions of this trade-off—what Bilbiie (2021) calls the *Catch-22*—have been put forward. In contrast to our model, they all rely on a specific design for either monetary or fiscal policy. Bilbiie (2021) shows that if monetary policy follows a Wicksellian price level targeting rule or fiscal policy follows a nominal bond rule, his tractable HANK model can simultaneously account for facts (i)-(iv).[37] Hagedorn et al. (2019) shows how introducing nominal government bonds and coupling it with a particular nominal bond supply rule can resolve the forward guidance puzzle in a quantitative HANK model (following the theoretical arguments in Hagedorn (2016) and Hagedorn (2018)). In contrast, we account for the four facts even in the case in which monetary policy follows a standard Taylor rule and absent any nominal bonds or specific fiscal rules.

Farhi and Werning (2019) also combine household heterogeneity with some form of bounded rationality, but focus entirely on resolving the forward-guidance puzzle. Our model accounts for a number of additional empirical facts, such as the transmission of monetary policy through indirect effects in a setting with unequal exposure of households to monetary policy and countercyclical income risk. We also consider a different form of bounded rationality, cognitive discounting, while Farhi and Werning (2019) focus on level-k thinking. Our setup is consistent with the empirical findings in Roth et al. (2021) who show that households adjust their interest-rate expectations only by about half of what the Fed announces, even when being told the Fed's intended interest-rate path.[38] In contrast to these papers, we consider supply shocks and show that the interaction of household heterogeneity and bounded rationality has qualitatively different implications for supply shocks than for forward guidance shocks.

Few other papers share the combination of nominal rigidities, household heterogeneity and some deviation from full information rational expectations (FIRE). Laibson et al. (2021) introduces *present bias* in a model of household heterogeneity but the model is set in partial equilibrium and they do not consider how the power of forward guidance or the stability

[37]Bilbiie (2021) proposes an additional resolution: a pure risk channel which can, in theory, break the co-movement of income risk and inequality. However, it requires a calibration which is at odds with the data.

[38]In an extension, we consider the case in which some households (financial markets, for example) fully incorporate the announced interest-rate paths into their expectations (see Section 4.4.3 where we discuss heterogeneous degrees of cognitive discounting) and show that our results remain robust in that scenario.

at the lower bound are affected by the presence of the two frictions. Auclert et al. (2020) incorporate sticky information into a HANK model to generate hump-shaped responses of macroeconomic variables to aggregate shocks while simultaneously matching intertemporal MPCs. Their paper, however, does not discuss the implications of the deviation from FIRE and heterogeneity for forward guidance or stability at the lower bound.[39]

Outline. The rest of the paper is structured as follows. We present our behavioral HANK model in Section 4.2. In Section 4.3, we consider a special calibration that allows us to solve the model in closed form and, thus, to build intuition for our results. In Section 4.4, we then move to a more standard calibration and show that all results remain robust in that case. We further discuss the role of heterogeneity in the behavioral bias for our results and non-rationality with respect to the idiosyncratic risk. We then use the quantitative model to study the policy implications of inflationary supply-side shocks in Section 4.5. Section 4.6 concludes.

4.2 Model

This section presents our model that incorporates household heterogeneity, cognitive discounting, and nominal rigidities. Our baseline model assumes sticky prices and flexible wages. We also redo our exercises with sticky wages (in Section 4.3.5 and 4.4.2)later on and show that our results do not depend on the specific nominal rigidities.

4.2.1 Households

Time is discrete and denoted by $t = 0, 1, 2, ...$ The economy is populated by a unit mass of households, indexed by $i \in [0, 1]$. Households obtain utility from (non-durable) consumption, $C_{i,t}$, and dis-utility from working $N_{i,t}$. Households discount future utility at rate $\beta_{i,t} \in (0, 1)$. We assume a standard CRRA utility function

$$\mathcal{U}(C_{i,t}, N_{i,t}) \equiv \begin{cases} \frac{C_{i,t}^{1-\gamma}}{1-\gamma} - \frac{N_{i,t}^{1+\varphi}}{1+\varphi}, & \text{if } \gamma \neq 1, \\ \log\left(C_{i,t}\right) - \frac{N_{i,t}^{1+\varphi}}{1+\varphi}, & \text{if } \gamma = 1, \end{cases} \tag{4.1}$$

where φ denotes the inverse Frisch elasticity and γ the relative risk aversion.

Household i faces the budget constraint

$$C_{i,t} + \frac{B_{i,t+1}}{1+r_t} = B_{i,t} + W_t z(e_{i,t}) N_{i,t} + D_t d(e_{i,t}) - \tau_t(e_{i,t}) \tag{4.2}$$

[39] Wiederholt (2015), Angeletos and Lian (2018), Andrade et al. (2019), Gabaix (2020) consider deviations from FIRE and Michaillat and Saez (2021) introduce wealth in the utility function (all in non-HANK setups) and show how to resolve the forward guidance puzzle. See, e.g., Broer et al. (2022) and Ilut and Valchev (2023) for recent contributions to how household heterogeneity and deviations from FIRE interact in settings abstracting from nominal rigidities.

and the borrowing constraint $B_{i,t+1} \geq \underline{B}$, where $\underline{B}$ denotes an exogenous borrowing limit, B denotes the household's bond holdings, r_t denotes the net real interest rate, W_t the real wage and $e_{i,t}$ the household's exogeneous idiosyncratic state that follows a Markov chain with time-invariant transition matrix $\mathcal{P}$. The process for $e_{i,t}$ is the same for all households and the mass of households in state e at any point in time equals the probability of being in that state in the stationary equilibrium, $p(e)$. Conditional on their exogenous idiosyncratic state, households have the idiosyncratic productivity $z(e_{i,t})$, they receive a share $d(e_{i,t})$ of total dividends D_t, and pay taxes $\tau_t(e_{i,t})$. We introduce taxes in such a way that they are non-distortionary in the sense that they do not show up in the household's first-order conditions. We also allow households' time discount factor to be a function of e, $\beta(e_{i,t})$.

Given their beliefs, households maximize their expected lifetime utility subject to their budget constraint (4.2) and the borrowing constraint. This yields the Euler equation

$$C_{i,t}^{-\gamma} \geq \beta(e_{i,t}) R_t \mathbb{E}_t^{BR} \left[C_{i,t+1}^{-\gamma} \right], \tag{4.3}$$

and the labor-leisure equation

$$N_{i,t}^{\varphi} = z(e_{i,t}) W_t C_{i,t}^{-\gamma}, \tag{4.4}$$

where $R_t \equiv 1 + r_t$ denotes the gross real interest rate. The Euler equation (4.3) holds with equality when the borrowing constraint does not bind, while it holds with strict inequality when the borrowing constraint binds. $\mathbb{E}_t^{BR}$ denotes the *boundedly-rational* expectations operator which we discuss next.

Bounded rationality. We assume that households are fully rational with respect to their idiosyncratic risk, but they cognitively discount the effects of aggregate shocks (we relax the assumption that households are rational with respect to their idiosyncratic risk in Section 4.4.4). To model cognitive discounting, we follow Gabaix (2020) but extend it to an economy with a whole distribution of households rather than focusing on a representative consumer.[40] Let X_t be a random variable (or vector of variables) and let us define $\bar{X}_t$ as some default value the agent may have in mind and let $\widetilde{X}_{t+1} \equiv X_{t+1} - \bar{X}_t$ denote the deviation from this default value.[41] The behavioral agent's expectation about X_{t+1} is then defined as

$$\mathbb{E}_t^{BR} \left[X_{t+1} \right] = \mathbb{E}_t^{BR} \left[\bar{X}_t + \widetilde{X}_{t+1} \right] \equiv \bar{X}_t + \bar{m} \mathbb{E}_t \left[\widetilde{X}_{t+1} \right], \tag{4.5}$$

[40] While Gabaix (2020) embeds bounded rationality in a NK model the basic idea of behavioral inattention (or sparsity) has been proposed by Gabaix earlier already (see Gabaix (2014) and Gabaix (2017)) and a handbook treatment of behavioral inattention is given in Gabaix (2019). We present a way how to microfound cognitive discounting as a noisy-signal extraction problem in Appendix D.4.7, but note, that the exact microfoundation or underlying behavioral friction which leads to underreaction is not crucial for the rest of our analysis. Angeletos and Lian (2017) show how forms of incomplete information can lead to observationally-equivalent expectations.

[41] Gabaix (2020) focuses on the case in which X_t denotes the state of the economy. He shows (Lemma 1 in Gabaix (2020)) that this form of cognitive discounting also applies to all other variables. Appendix D.4.6 derives our results following the approach in Gabaix (2020). The results remain exactly the same.

where $\mathbb{E}_t[\cdot]$ is the rational expectations operator and $\bar{m} \in [0,1]$ is the cognitive discounting parameter. A higher $\bar{m}$ denotes a smaller deviation from rational expectations and rational expectations are captured by $\bar{m} = 1$. Our setup therefore nests the rational expectations model as a special case.

When $\bar{m} < 1$, the behavioral agent anchors her expectations to the default value and cognitively discounts expected future deviations from this default value. Given that households are perfectly rational with respect to their idiosyncratic risk and only cognitively discount the implications of aggregate shocks, we assume that the default value $\bar{X}_t$ is given by the variable's stationary equilibrium counterpart. Thus, when there is no aggregate shock and the economy is in the stationary equilibrium, $\widetilde{X}_{t+1} = 0$, households are fully rational.

To see how cognitive discounting matters in our model, note that the only forward-looking equation in the household block is the Euler equation (4.3). Let $\bar{C}_{i,t} \equiv C(e_{i,t}, B_{i,t}, \bar{Z})$ denote consumption of household i in period t with exogeneous idiosyncratic state $e_{i,t}$ and asset holdings $B_{i,t}$ when all aggregate variables are in steady state, indicated by $\bar{Z}$. Here, Z potentially denotes a whole matrix of aggregate variables, including, for example, news shocks (i.e., forward guidance shocks). In other words, $\bar{C}_{i,t}$ denotes consumption of household i with exogeneous state $e_{i,t}$ and asset holdings $B_{i,t}$ in the stationary equilibrium, and thus, the household's default (or anchor) value of consumption. In case an aggregate shock occurs, $Z_t \neq \bar{Z}$, consumption is denoted by $C_{i,t} = C(e_{i,t}, B_{i,t}, Z_t)$. We can then write the Euler equation with bounded rationality (BR) in terms of the rational expectations operator $\mathbb{E}_t[\cdot]$ as

$$\begin{aligned} C_{i,t}^{-\gamma} &\geq \beta(e_{i,t}) R_t \mathbb{E}_t^{BR}\left[C_{i,t+1}^{-\gamma}\right] \\ &= \beta(e_{i,t}) R_t \mathbb{E}_t^{BR}\left[\bar{C}_{i,t+1}^{-\gamma} + \left(C_{i,t+1}^{-\gamma} - \bar{C}_{i,t+1}^{-\gamma}\right)\right] \\ &= \beta(e_{i,t}) R_t \mathbb{E}_t\left[\bar{C}_{i,t+1}^{-\gamma} + \bar{m}\left(C_{i,t+1}^{-\gamma} - \bar{C}_{i,t+1}^{-\gamma}\right)\right], \end{aligned} \tag{4.6}$$

where the rational expectations operator $\mathbb{E}_t[\cdot]$ denotes the expectations that a fully rational household would have in the behavioral economy.

Equation (4.6) illustrates that when households form expectations about their marginal utility in the next period, their expectations about the marginal utilities associated with each possible individual state are anchored to the marginal utilities associated with these states in stationary equilibrium. Thus, the household's default value of her future marginal utility is a whole distribution of marginal utilities, depending on her individual state $(e_{i,t}, B_{i,t})$.

Underreaction in the data. Given $\bar{m} < 1$, expectations underreact to aggregate news about the future compared to the rational expectations case, that is, they do not fully incorporate aggregate news into their expectations. We now show that households indeed show patterns of underreaction in the data. We follow Coibion and Gorodnichenko (2015) and regress forecast errors on forecast revisions as follows

$$\underbrace{x_{t+4} - \mathbb{E}_t^{e,BR} x_{t+4}}_{\text{Forecast errors}} = c^e + b^{e,CG} \underbrace{\left(\mathbb{E}_t^{e,BR} x_{t+4} - \mathbb{E}_{t-1}^{e,BR} x_{t+3}\right)}_{\text{Forecast revision}} + \epsilon_t^e, \tag{4.7}$$

and we do so for different income groups, indexed by e. As we show in Appendix D.2, $b^{e,CG} > 0$ is consistent with underreaction and the corresponding cognitive discounting parameter can be obtained from

$$\bar{m}^e = \left(\frac{1}{1 + b^{e,CG}}\right)^{1/4}. \tag{4.8}$$

As we focus on bounded rationality with respect to aggregate shocks, we consider expectations about aggregate variables, namely, unemployment changes, the unemployment level, and inflation which we obtain from the Survey of Consumers from the University of Michigan. We split households into three groups based on their income. The bottom and top income groups each contain the 25% households with the lowest and highest income, respectively, and the remaining 50% are assigned to the middle income group. As the expectations in the forecast revisions in equation (4.7) are about the variable at different points in time (due to data limitations), we instrument forecast revisions by the *main business cycle shock* obtained from Angeletos et al. (2020).

We find that in all cases $\widehat{b}^{e,CG}$ is positive, suggesting that households of all income groups tend to underreact, consistent with our assumption of $\bar{m} < 1$ (Table D.1 in the Appendix provides the details). Using equation (4.8) we obtain estimates of $\bar{m}^e$ equal to 0.57, 0.59 and 0.64 for the bottom 25%, the middle 50% and the top 25%, respectively for the estimates from the IV regressions when focusing on expected unemployment changes. When we consider unemployment levels rather than changes, the estimated $\bar{m}^e$ equal 0.86, 0.87 and 0.88. If we consider inflation expectations instead of unemployment expectations, we obtain estimated cognitive discounting parameters of 0.70, 0.75 and 0.78 for the bottom 25%, the middle 50% and the top 25%, respectively.[42]

There are two take-aways from this empirical exercise: first, households of all income groups underreact in their expectations. Second, the estimated cognitive discounting parameters tend to be between 0.6 and 0.85, consistent with values used in Gabaix (2020).

Consistent with our findings, Kučinskas and Peters (2022) and Born et al. (2022) find that professional forecasters and firms, respectively, underreact to aggregate shocks. However, they also find evidence of *overreaction* (as in Bordalo et al. (2020)) to idiosyncratic shocks. We discuss this case where households underreact to aggregate shocks but overreact to idiosyncratic shocks in Section 4.4.4.

[42] Estimates using OLS rather than IV are similar (see Appendix D.2).

4.2.2 Firms

We assume a standard New Keynesian firm side with sticky prices and where firms have rational expectations (the case with flexible prices and sticky wages is discussed in Sections 4.3.5 and 4.4.2, and the case with boundedly-rational firms in 4.5 and Appendix D.4.5). All households consume the same aggregate basket of individual goods, $j \in [0,1]$, $C_t = \left(\int_0^1 C_t(j)^{\frac{\epsilon-1}{\epsilon}} dj\right)^{\frac{\epsilon}{\epsilon-1}}$, where $\epsilon > 1$ is the elasticity of substitution between the individual goods. Each firm faces demand $C_t(j) = \left(\frac{P_t(j)}{P_t}\right)^{-\epsilon} C_t$, where $P_t(j)/P_t$ denotes the individual price relative to the aggregate price index, $P_t^{1-\epsilon} = \int_0^1 P_t(j)^{1-\epsilon} dj$, and produces with the linear technology $Y_t(j) = N_t(j)$. Firms can only update their prices infrequently, as in Calvo (1983) and Yun (1996). The real marginal cost is given by W_t. We assume that the government pays a constant subsidy τ^S on revenues to induce marginal cost pricing in the steady state. This subsidy is financed by a lump-sum tax on firms T_t^F. Hence, the profit function is $D_t(j) = (1+\tau^S)[P_t(j)/P_t]Y_t(j) - W_t N_t(j) - T_t^F$. Total profits are then $D_t = Y_t - W_t N_t$ and are zero in steady state.

4.2.3 Government

The government consists of a fiscal authority and a monetary authority. The fiscal authority faces the budget constraint

$$\frac{B_{t+1}^G}{R_t} + T_t = B_t^G,$$

where B^G denotes the bonds issued by the government and T_t denotes tax income. We abstract from government spending. Taxes follow a simple debt feedback rule

$$T_t - \bar{T} = \vartheta \frac{B_{t+1}^G - \bar{B}^G}{\bar{Y}}, \tag{4.9}$$

where $\bar{T}$, $\bar{B}^G$ and $\bar{Y}$ denote the respective steady state values. Further, fiscal policy induces the optimal steady state subsidy financed by lump-sum taxation of firms.

In most of the analysis, we assume that monetary policy either sets the nominal interest rate i_t following a standard (linearized) Taylor rule

$$\hat{i}_t = \phi \pi_t + \epsilon_t^{MP}, \tag{4.10}$$

or a *real rate rule*

$$r_t = \bar{r} + \epsilon_t^{MP}, \tag{4.11}$$

with ϵ_t^{MP} being a monetary policy shock, π_t denoting inflation, $\bar{r}$ the steady-state real interest rate, and where variables with a "^" denote log deviations from the variables' respective steady state values. The parameter ϕ captures how strongly monetary policy responds to inflation. For now, monetary policy shocks are the only source of aggregate uncertainty.

Equilibrium definition. Given an initial price level P_{-1}, initial government debt level B_0^G, and an initial distribution of agents $\Psi_0(B_0, e_0)$, a general equilibrium is a path for prices $\{P_t, W_t, \pi_t, r_t, i_t\}$, aggregates $\{Y_t, C_t, N_t, B_{t+1}^G, T_t, D_t\}$, individual allocation rules $\{C_t(B_t, e_t), B_{t+1}(B_t, e_t)\}$ and joint distributions of agents $\Psi_t(B_t, e_t)$ such that households optimize (given their beliefs), all firms optimize, monetary and fiscal policy follow their rules, and the goods and bond markets clear:

$$\sum_e p(e) \int C_t(B_t, e_t)\, \Psi_t(B_t, e_t) = Y_t$$
$$\sum_e p(e) \int B_{t+1}(B_t, e_t)\, \Psi_t(B_t, e_t) = B_{t+1}^G.$$

4.3 Analytical Results

To understand how household heterogeneity and cognitive discounting interact, we now calibrate the model such that we can solve the model in closed form. We refer to this specific calibration as *tractable* behavioral HANK model as it nests the tractable rational HANK model of Bilbiie (2020) and Bilbiie (2021).

4.3.1 A Calibration towards a Closed-Form Solution

Solving the model in closed form requires specific functional forms for $\beta(e)$, $z(e)$, $d(e)$, and $\tau(e)$, for the stochastic process of e, as well as $B_t^G = \underline{B} = 0$ for all t. Starting with the process of e, we for now assume that there are only two states, $e \in \{U, H\}$, and denote a household's probability to remain in her current state $p(e_{t+1} = U | e_t = U) = s$ and $p(e_{t+1} = H | e_t = H) = h$. Consequently, $\lambda = \frac{1-s}{2-s-h}$ is the time-constant share of households being in state H. We then assume that $\beta(H) < \beta(U)$ such that the Euler equation (4.6) always holds with equality for households being in state U, while it always holds with inequality for households being in state H. In other words, H households are always $H$and-to-Mouth, while U households are always Unconstrained. In addition, we assume that $z(e) = 1$ and $\tau(e) = 0$ for both states and $d(H) = \frac{\mu^D}{\lambda}$ and $d(U) = \frac{1-\mu^D}{1-\lambda}$. This leaves two sources of income heterogeneity, namely, different labor supply and different profit shares.

The assumption that $B_t^G = \underline{B} = 0$ for all t means that the government does not issue any bonds and households cannot borrow. It follows that households cannot save in equilibrium and therefore, all H households are identical and all U households are identical, independent of how long they have been in state U or in state H. We can thus solve the model as if there were two representative households, a Hand-to-Mouth and an Unconstrained household. Hence, in this section, we will use superscripts H and U to indicate the two representative households.

As profits are zero in steady state due to the subsidy induced by fiscal policy, it follows that households are identical in steady state, $C^H = C^U = C$. In the log-linear dynamics around this steady state, profits vary inversely with the real wage, $\widehat{d}_t = -\widehat{w}_t$. We allow for steady state inequality in Appendix D.4 and show that our results are not driven by this assumption.

4.3.2 Log-Linearized Dynamics

We now focus on the log-linearized dynamics around the full-insurance, zero-liquidity steady state. The first key equilibrium equation is the consumption of the hand-to-mouth households written as a function of total output

$$\widehat{c}_t^H = \chi \widehat{y}_t, \tag{4.12}$$

with

$$\chi \equiv 1 + \varphi\left(1 - \frac{\mu^D}{\lambda}\right) \tag{4.13}$$

measuring the cyclicality of the H household's consumption (see appendix D.1.1). Patterson (2023) documents that households with higher MPCs tend to be more exposed to aggregate income fluctuations induced by monetary policy or other demand shocks—fact (ii) in the introduction. We can account for fact (ii) by setting $\chi > 1$. Similarly, Auclert (2019) finds that poorer households tend to exhibit higher MPCs. Together with the findings in Coibion et al. (2017) and Hintermaier and Koeniger (2019) that poorer households' income is on average more exposed to monetary policy shocks, this also implies $\chi > 1$. For given φ, this requires $\mu^D < \lambda$.

Why does $\mu^D < \lambda$ imply that the consumption of hand-to-mouth households moves more than one-for-one with aggregate output after a monetary policy shock? Consider an expansionary monetary policy shock, i.e., an unexpected decrease in the interest rate. Unconstrained households want to consume more and save less, leading to an increase in demand. Firms then increase their labor demand, leading to an increase in wages. Due to the assumption of sticky prices and flexible wages, profits in the New Keynesian model decrease ($\widehat{d}_t = -\widehat{w}_t$). In the representative agent model, the representative agent both incurs the increase in wages and the decrease in profits coming from firms. With household heterogeneity, however, this is not necessarily the case. If $d(H) < 1$, which is the case when $\mu^D < \lambda$, the decrease in profits affects the income of H households less than one-for-one while the increase in the real wage affects their income one-for-one. Thus, the total income of H households increases more than one-for-one with aggregate income.

Combining equation (4.12) with the goods market clearing condition yields

$$\widehat{c}_t^U = \frac{1 - \lambda\chi}{1 - \lambda}\widehat{y}_t, \tag{4.14}$$

which implies that consumption inequality is given by:[43]

$$\widehat{c}_t^U - \widehat{c}_t^H = \frac{1-\chi}{1-\lambda}\widehat{y}_t. \tag{4.15}$$

Thus, if $\chi > 1$, inequality is countercyclical as it varies negatively with total output, i.e., inequality increases in recessions and decreases in booms. In line with the empirical evidence on the covariance between MPCs and income exposure, the data also points towards $\chi > 1$ when looking at the cyclicality of inequality, conditional on monetary policy: Coibion et al. (2017), Mumtaz and Theophilopoulou (2017), Ampudia et al. (2018) and Samarina and Nguyen (2019) all provide evidence of countercyclical inequality conditional on monetary policy shocks.

The second key equilibrium equation is the log-linearized bond Euler equation of U households:

$$\widehat{c}_t^U = s\mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^U\right] + (1-s)\mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^H\right] - \frac{1}{\gamma}\left(\widehat{i}_t - \mathbb{E}_t^{BR}\pi_{t+1}\right). \tag{4.16}$$

For the case without idiosyncratic risk, i.e., for $s = 1$, equation (4.16) boils down to a standard Euler equation under bounded rationality. For $s \in [0, 1)$, however, the household takes into account that she might be hit by an idiosyncratic shock and self-insures against becoming hand-to-mouth next period. How strongly this precautionary savings motive affects the household's consumption away from the stationary equilibrium will depend on the household's degree of bounded rationality. We will, following the assumption in Gabaix (2020), often focus on the case in which households are rational with respect to today's real rate, i.e., we replace $\mathbb{E}_t^{BR}\pi_{t+1}$ with $\mathbb{E}_t\pi_{t+1}$ in equation (4.16). We show in Appendix D.4 that our results go through with boundedly-rational expectations of today's real rate.

Supply side. For simplicity and to get a clear understanding of the mechanisms driving our results, we focus on a static Phillips curve in this section:

$$\pi_t = \kappa\widehat{y}_t, \tag{4.17}$$

where $\kappa \geq 0$ captures the slope of the Phillips curve. Such a static Phillips curve arises if we assume that firms are either completely myopic or if they face Rotemberg-style price adjustment costs relative to yesterday's market average price index, instead of their own price (see Bilbiie (2021)). In Appendix D.4.5, we show that a forward-looking Phillips Curve (rational or behavioral) does not qualitatively affect our results.

[43]We denote the case in which unconstrained households consume relatively more than hand-to-mouth households as higher inequality, even though they consume the same amount in steady state. As we move away from the tractable model in Sections 4.4 and 4.5, households' consumption levels will differ in the stationary equilibrium.

4.3.3 The Closed-Form Solution

Our tractable behavioral HANK model can be summarized by three equations: a Phillips curve, representing the aggregate supply side captured by equation (4.17), a rule for monetary policy (equation (4.10) or (4.11)), which together with the third equation—the aggregate IS equation—determines aggregate demand. To obtain the aggregate IS equation, we combine the hand-to-mouth households' consumption (4.12) with the consumption of unconstrained households (4.14) and their Euler equation (4.16) (see appendix D.1 for all the derivations).

Proposition 5 *The aggregate IS equation is given by*

$$\widehat{y}_t = \psi_f \mathbb{E}_t \widehat{y}_{t+1} - \psi_c \frac{1}{\gamma} \left(\widehat{i}_t - \mathbb{E}_t \pi_{t+1} \right), \tag{4.18}$$

where

$$\psi_f \equiv \bar{m}\delta = \bar{m} \left[1 + (\chi - 1) \frac{1-s}{1-\lambda\chi} \right] \quad \textit{and} \quad \psi_c \equiv \frac{1-\lambda}{1-\lambda\chi}.$$

Compared to the rational representative agent model, two new coefficients show up: ψ_c and ψ_f. ψ_c governs the sensitivity of today's output with respect to the contemporaneous real interest rate. ψ_c is shaped by household heterogeneity, in particular by the share of H households λ and their income exposure χ. As the H households' incomes are more exposed to aggregate income ($\chi > 1$), $\psi_c > 1$ which renders current output more sensitive to changes in the contemporaneous real interest rate due to general equilibrium forces, as we show later.

The second new coefficient in the behavioral HANK IS equation (4.18), ψ_f, captures the sensitivity of today's output with respect to changes in expected future output. ψ_f is shaped by household heterogeneity and the behavioral friction as it depends on the precautionary-savings motive, captured by δ, *and* the degree of bounded rationality of households as well as the interaction of these two. Given that $\chi > 1$, unconstrained households take into account that they will be more exposed to aggregate income fluctuations in case they become hand-to-mouth. Thus, income risk is countercyclical, which manifests itself in $\delta > 1$ (consistent with the empirical evidence, e.g., in Storesletten et al. (2004) or Guvenen et al. (2014)). Countercyclical risk induces compounding in the Euler equation and, thus, competes with the empirically observed underreaction of aggregate expectations ($\bar{m} < 1$) which induces discounting in the Euler equation. We see in the following sections that even for a small degree of bounded rationality—much smaller than the empirics suggest—that discounting through bounded rationality dominates the compounding through countercyclical income risk. Hence, in the behavioral HANK model it holds that $\psi_f < 1$ which makes the economy less sensitive to expectations and news about the future.

Equation (4.18) nests IS equations of three classes of models in the literature: first, the representative-agent rational expectations (RANK) model which can be obtained by setting

$\bar{m} = 1$ and assuming only one state $e = U$ which would imply $\psi_f = \psi_c = 1$ (see Galí (2015), Woodford (2003)).[44] Second, representative agent models deviating from full-information rational expectations when assuming one state and $\bar{m} \in (0, 1)$ which results in $\psi_c = 1$ and $\psi_f < 1$ as, for example, in Gabaix (2019), Angeletos and Lian (2018) and Woodford (2019). And thirdly, TANK and tractable HANK models as e.g., in Bilbiie (2008), Bilbiie (2021), McKay et al. (2017), or Debortoli and Galí (2018) with again two states but $\bar{m} = 1$ which implies $\psi_f = \delta$. Nesting these models enables us to clearly illustrate why our model is able to account for fact (i) - (iv) simultaneously while these other models cannot.

Calibration. Given that this stylized version of our model is a tractable HANK model with two agents, we calibrate it using standard parameters in the literature on tractable HANK models (see, e.g., Bilbiie (2020), Bilbiie (2021)).[45] That said, we show in Appendix D.4.1 that our results are robust to a wide range of parameters. We set the share of H agents to one third, $\lambda = 0.33$, and μ^D such that $\chi = 1.35$ which implies $\psi_c = 1.2$. We set $\chi > 1$ to capture that high-MPC households' incomes are relatively more sensitive to aggregate fluctuations induced by monetary policy, in line with the findings in Patterson (2023), Coibion et al. (2017) and Auclert (2019). We set the probability of a U household to become hand-to-mouth next period to 5.4%, i.e., $s = 0.946$ (this corresponds to $s = 0.8$ in annual terms). We focus on log utility, $\gamma = 1$, set $\beta(U) = 0.99$, and the slope of the Phillips Curve to $\kappa = 0.02$, as in Bilbiie et al. (2022). The cognitive discounting parameter, $\bar{m}$ is set to 0.85, as explained in Section 4.2. Note, that even when we vary certain parameters, we keep $\lambda < \chi^{-1}$.

4.3.4 Monetary Policy

We now show how the behavioral HANK model generates amplification of contemporaneous monetary policy through indirect effects while resolving the forward guidance puzzle at the same time. Additionally, we discuss determinacy conditions and show that the model remains stable at the effective lower bound.

General equilibrium amplification and forward guidance. We start by showing how the behavioral HANK model generates amplification of current monetary policy through indirect general equilibrium effects while simultaneously ruling out the forward guidance puzzle. The forward guidance puzzle states that announcements about future changes in the interest rate affect output today as strong (or even stronger) than contemporaneous changes in the interest rate.[46] Such strong effects of future interest rate changes, however, seem puzzling and are not supported by the data (Del Negro et al. (2015), Roth et al. (2021)).

[44]Only one state implies that χ vanishes from the model and $\lambda = 0$ and $s = 1$.

[45]In the next section, we then show how we can use our quantitative model to directly match micro evidence from Patterson (2023) on the unequal income exposure of households.

[46]Detailed analyses of the forward guidance puzzle in RANK are provided by McKay et al. (2016) and Del Negro et al. (2015).

Let us consider two different monetary policy experiments: (i) a contemporaneous monetary policy shock, i.e., a surprise decrease in the real interest rate today, and (ii) a forward guidance shock, i.e., a news shock today about a decrease in the real interest rate k periods in the future. The monetary authority keeps the real interest rate at its steady state value in all other periods. We focus on real rate changes as this is the set up that McKay et al. (2016) focus on and Farhi and Werning (2019) focus on the case with fully-rigid prices, such that nominal rate changes translate one-for-one to real rate changes. However, all our results are robust when focusing on nominal rate changes and are presented in Appendix D.4.2

Proposition 6 *In the behavioral HANK model, there is amplification of contemporaneous monetary policy relative to RANK if and only if*

$$\psi_c > 1 \Leftrightarrow \chi > 1, \tag{4.19}$$

and the forward guidance puzzle is ruled out if

$$\psi_f < 1. \tag{4.20}$$

Let us first focus on equation (4.19) which tells us that the behavioral HANK model generates amplification of contemporaneous monetary policy with respect to RANK whenever $\chi > 1$, that is, when high-MPC households' consumption is relatively more sensitive to aggregate income fluctuations.

After a decrease in the interest rate, wages increase and profits decline. As H agents receive a relatively smaller share of profits but fully benefit from the increase in wages, their income increases more than one-to-one with aggregate income. As they consume their income immediately, the initial effect on total output increases. The unconstrained households, on the other hand, experience a smaller increase in their income due to the fall in their profit income. As a result, $\psi_c > 1$ and the increase in output is amplified through these general equilibrium effects. To see the importance of general equilibrium or indirect effects, the following Lemma disentangles the direct and indirect effects.

Lemma 2 *The consumption function in the behavioral HANK model is given by*

$$\widehat{c}_t = [1 - \beta(1 - \lambda\chi)]\,\widehat{y}_t - \frac{(1-\lambda)\beta}{\gamma}\widehat{r}_t + \beta\bar{m}\delta(1-\lambda\chi)\mathbb{E}_t\widehat{c}_{t+1}. \tag{4.21}$$

Let ρ denote the exogenous persistence and define the indirect effects as the change in total consumption due to the change in total income but for fixed real rates. The share of indirect effects, Ξ^{GE}, out of the total effect is then given by

$$\Xi^{GE} = \frac{1-\beta(1-\lambda\chi)}{1-\beta\bar{m}\delta\rho(1-\lambda\chi)}.$$

Given our calibration and assuming an AR(1) monetary policy shock with a persistence of 0.6, indirect effects account for about 63%, consistent with larger quantitative models as for example in Kaplan et al. (2018) and thus, the model accounts for fact (i).[47] Holm et al. (2021) state that the overall importance of indirect effects they find in the data is comparable to those in Kaplan et al. (2018), with the difference that these effects unfold after some time, whereas direct effects are more important on impact. Because in our stylized model the response to a monetary policy shock peaks on impact, indirect effects are important right away. Slacalek et al. (2020) provide further evidence that indirect effects are strong drivers of aggregate consumption in response to monetary policy shocks. For comparison, the representative agent model generates an indirect share of $\Xi^{GE} = \frac{1-\beta}{1-\beta\bar{m}\rho}$, which, given our calibration, amounts to about 2%.

Turning to forward guidance, equation (4.20) in Proposition 6 tells us that the forward guidance puzzle is ruled out if $\psi_f < 1$. What determines whether this condition holds or not? First, note that as in the discussion of contemporaneous monetary policy, with $\chi > 1$ the income of H agents moves more than one for one with aggregate income. In this case, unconstrained households who self-insure against becoming hand-to-mouth in the future want less insurance when they expect a decrease in the interest rate because if they become hand-to-mouth they would benefit more from the increase in aggregate income. Hence, after a forward guidance shock, unconstrained households decrease their precautionary savings which compounds the increase in output today ($\delta > 1$). Yet, as households are boundedly rational, they cognitively discount these effects taking place in the future. Importantly, unconstrained households cognitively discount both the usual consumption-smoothing response due to the future increase in consumption as well as the general equilibrium implications for their precautionary savings, thereby decreasing the effects of the forward guidance shock on today's consumption. Thus, the model not only accounts for facts (i) and (ii) but simultaneously accounts for fact (iii).

This last part clearly illustrates the main interaction of bounded rationality and household heterogeneity that enables the behavioral HANK model to resolve the forward guidance puzzle while simultaneously generating amplification through indirect effects. Households fully understand their idiosyncratic risk of switching their type as well as the implications of switching type in case there are no aggregate shocks, i.e., in the steady state. If the monetary authority makes an unexpected announcement about its future policy, however, behavioral households do not fully incorporate the effects of this policy on their own income risk and thus, their precautionary savings. Already a small underreaction of the behavioral households is enough to resolve the forward guidance puzzle. Given our calibration there is no forward guidance puzzle in the behavioral HANK model as long as $\bar{m} < 0.966$ which is above the upper bounds for empirical estimates (see Section 4.2).[48] Figure D.4 in Appendix

[47] We write β for $\beta(U)$ for notational simplicity and because $\beta(H)$ does not affect any of our results (as long as it is low enough such that the borrowing constraint always binds for H households).

[48] A related paradox in the rational model is that as the persistence of the shock increases, the effects

D.4.1 shows that the solution of the forward guidance in our model is very robust with respect to changes in the heterogeneity parameters.

We now compare the behavioral HANK model to its rational counterpart to show how the behavioral HANK model overcomes a shortcoming inherent in the rational HANK model – the *Catch-22* (Bilbiie (2021); see also Werning (2015)). The Catch-22 describes the tension that the rational HANK model can either generate amplification of contemporaneous monetary policy *or* solve the forward guidance puzzle. To see this, note that with $\bar{m} = 1$ the forward guidance puzzle is resolved when $\delta < 1$ which requires $\chi < 1$, as otherwise $\delta > 1$. Assuming $\chi < 1$, however, leads to *dampening* of contemporaneous monetary policy instead of amplification. We graphically illustrate the Catch-22 of the rational model and its resolution in the behavioral HANK model in Appendix D.3. Note that also rational TANK models (thus, turning off type switching) or the behavioral RANK model would not deliver amplification and resolve the forward guidance puzzle simultaneously. TANK models would face the same issues as the rational RANK model in the sense that they cannot solve the forward guidance puzzle while bounded rationality in a RANK model does not deliver initial amplification through indirect effects.

Stability at the Effective Lower Bound

In this section, we revisit the determinacy conditions in the behavioral HANK model and discuss the implications for the stability at the effective lower bound constraint on nominal interest rates. We therefore focus on the case where monetary policy follows the Taylor rule (4.10) (we discuss more general Taylor rules in Appendix D.1.4). To derive these results, it is sometimes convenient to combine the IS equation (4.18) with the static Phillips Curve (4.17) and the Taylor rule (4.10) so that we can represent the model in a single first-order difference equation:

$$\widehat{y}_t = \frac{\psi_f + \psi_c \frac{\kappa}{\gamma}}{1 + \psi_c \phi \frac{\kappa}{\gamma}} \mathbb{E}_t \widehat{y}_{t+1} - \frac{\psi_c \frac{1}{\gamma}}{1 + \psi_c \phi \frac{\kappa}{\gamma}} \varepsilon_t^{MP}. \tag{4.22}$$

According to the Taylor principle, monetary policy needs to respond sufficiently strongly to inflation in order to guarantee a determinate equilibrium. In the rational RANK model the Taylor principle is given by $\phi > 1$, where ϕ is the inflation-response coefficient in the Taylor rule (4.10). We now derive a similar determinacy condition in the behavioral HANK model and show that both household heterogeneity and bounded rationality affect this condition. The following proposition provides the behavioral HANK Taylor principle.[49]

Proposition 7 *The behavioral HANK model has a determinate, locally unique equilibrium*

become unboundedly large and as the persistence approaches unity, an exogenous increase in the nominal interest rate becomes expansionary. The behavioral HANK model, on the other hand, does not suffer from this. We elaborate these points in more detail in Appendix D.4.4.

[49] We focus on local determinacy and bounded equilibria.

if and only if:

$$\phi > \phi^* = 1 + \frac{\bar{m}\delta - 1}{\frac{\kappa}{\gamma}\frac{1-\lambda}{1-\lambda\chi}}. \tag{4.23}$$

We obtain Proposition 7 directly from the difference equation (4.22). For determinacy, we need that the coefficient in front of $\mathbb{E}_t\widehat{y}_{t+1}$ is smaller than 1 (the eigenvalues associated with any exogenous variables are assumed to be $\rho < 1$, and are thus stable). Solving this condition for ϕ yields Propositon 7. Appendix D.1.4 outlines the details and extends the result to more general Taylor rules.

To understand the condition in Proposition 7, consider first $\bar{m} = 1$ and, thus, focus solely on the role of household heterogeneity. With $\chi > 1$, it follows that $\phi^* > 1$ and, hence, the threshold is higher than the RANK Taylor principle states. This insufficiency of the Taylor principle in rational HANK models has been shown by Bilbiie (2021) and in a similar way by Ravn and Sterk (2021) and Acharya and Dogra (2020). As a future aggregate sunspot increases the income of households in state H disproportionately, unconstrained households cut back on precautionary savings today which further increases output today. This calls for a stronger response of the central bank to not let the sunspot become self-fulfilling.

On the other hand, bounded rationality $\bar{m} < 1$ relaxes the condition as unconstrained households now cognitively discount both the future aggregate sunspot as well as its implications for their idiosyncratic risk. A smaller response of the central bank is needed in order to prevent the sunspot to become self-fulfilling. Given our calibration the cutoff value for $\bar{m}$ to restore the RANK Taylor principle in the behavioral HANK model is 0.966. What is more, given our baseline choice of $\bar{m} = 0.85$, we obtain $\phi^* < 0$. Thus, in our tractable behavioral HANK model it is not necessary that monetary policy responds to inflation at all as the economy features a stable unique equilibrium even under an interest rate peg.

Stability at the effective lower bound. Related to the indeterminacy issues under a peg the traditional New Keynesian model struggles to explain how the economy can remain stable when the effective lower bound (ELB) on nominal interest rates is binding for an extended period of time, as observed in many advanced economies over recent decades (see, e.g., Debortoli et al. (2020) and Cochrane (2018)). If the ELB binds for a sufficiently long time, RANK predicts unreasonably large recessions and, in the limit case in which the ELB binds forever, even indeterminacy.[50] Similar to the forward guidance puzzle, this is even more severe in rational HANK models.

We now show that the behavioral HANK model resolves these issues and thus accounts for fact (iv). To this end, let us add a *natural rate shock* (i.e., a demand shock) $\widehat{r}_t^n$ to the IS

[50]A forever binding ELB basically implies that the Taylor coefficient is equal to zero and, thus, the nominal rate is pegged at the lower bound, thereby violating the Taylor principle. Note, that this statement also extends to models featuring more elaborate monetary policy rules including Taylor rules responding to output or also the Wicksellian price-level targeting rule, as they all collapse to a constant nominal rate in a world of an ever-binding ELB.

equation:

$$\widehat{y}_t = \psi_f \mathbb{E}_t \widehat{y}_{t+1} - \psi_c \left(\widehat{i}_t - \mathbb{E}_t \pi_{t+1} - \widehat{r}_t^n \right).$$

We assume that in period t the natural rate decreases to a value $\widetilde{r}^n$ that is sufficiently negative such that the natural rate in levels is below the ELB. The natural rate stays at $\widetilde{r}^n$ for $k \geq 0$ periods and after k periods the economy returns immediately back to steady state. Agents correctly anticipate the length of the binding ELB. Iterating the IS equation forward, it follows that output in period t is given by

$$\widehat{y}_t = -\frac{1}{\gamma} \psi_c \underbrace{\left(\widehat{i}_{ELB} - \widetilde{r}^n \right)}_{>0} \sum_{j=0}^{k} \left(\psi_f + \frac{\kappa}{\gamma} \psi_c \right)^j, \tag{4.24}$$

where the term $\left(\widehat{i}_{ELB} - \widetilde{r}^n \right) > 0$ captures the shortfall of the policy response due to the binding ELB. Under rational expectations, we have that $\psi_f > 1$ (and $\frac{\kappa}{\gamma} \psi_c > 0$), meaning that output implodes as $k \to \infty$. The same is true in the rational RANK model which is captured by $\psi_f = \psi_c = 1$. In the behavioral HANK model, however, this is not the case. As long as $\psi_f + \frac{\kappa}{\gamma} \psi_c < 1$ the output response in t is bounded even as $k \to \infty$. It follows that $\bar{m} < 0.94$ is enough to rule out unboundedly-severe recessions at the ELB even if the ELB is expected to persist forever. We graphically illustrate in Appendix D.3 that the behavioral HANK model remains stable also for long spells of the ELB in which output in the rational models collapses.

4.3.5 Sticky Wages

So far, we have assumed that prices are sticky and wages are fully flexible. This assumption, however, is not crucial for our aggregate results. The key difference in our context is the underlying mechanism for the unequal exposure of households.

To highlight this, we now consider wages to be sticky and prices to be fully flexible. Given that prices are fully flexible, we also abstract from monopolistic competition of firms, that is, prices are set to marginal costs. From the aggregate production function, $Y_t = N_t$, it follows that with flexible prices the aggregate price index P_t equals the nominal wage, such that that the real wage W_t is constant and equal to 1 (Auclert et al. (ming)). Further, a labor union allocates hours to workers and we assume that all households work the same amount in the steady state. If there is an aggregate shock, however, and hours deviate from their steady state value, $\widehat{n}_t \neq 0$, the labor union allocates these hours as follows:

$$\widehat{n}_t^H = \zeta \widehat{n}_t,$$

with ζ capturing the H households' sensitivity of hours worked to changes in total hours worked. Absent profits, taxes, and transfers, this allocation rule is the only source of income heterogeneity and, thus, ζ is a sufficient statistic of households' income exposure to monetary policy. Fact (ii)—that households with higher MPCs are more strongly exposed to monetary policy shocks—implies $\zeta > 1$. It directly follows that $\widehat{c}_t^H = \zeta \widehat{y}_t$ from the H households' budget constraint and the production function. Market clearing then yields $\widehat{c}_t^U = \frac{1-\lambda\zeta}{1-\lambda}\widehat{y}_t$.

Using these expressions for $\widehat{c}^H$ and $\widehat{c}^U$ in the unconstrained household's Euler equation yields

$$\frac{1-\lambda\zeta}{1-\lambda}\widehat{y}_t = s\bar{m}\frac{1-\lambda\zeta}{1-\lambda}\mathbb{E}_t\widehat{y}_{t+1} + (1-s)\bar{m}\zeta\mathbb{E}_t\widehat{y}_{t+1} - \frac{1}{\gamma}\widehat{r}_t,$$

which is exactly the same IS equation as the one in Proposition 5 when setting $\zeta = \chi$ (see Appendix D.1 for the algebra). Thus, the effects of conventional monetary policy shocks as well as the effects of forward guidance shocks on total output are exactly the same as in our baseline model with sticky prices and flexible wages. But instead of relying on the countercyclicality of profits, the model with sticky wages and flexible prices relies on the labor union's allocation rule of hours worked outside of the steady state to match fact (ii).

4.4 Quantitative Results

In this section, we relax the specific calibration choices that we use to solve the model in closed-form and show that all our results carry over. To this end, we build on a standard calibration in the HANK literature which implies that the model features a non-degenerate wealth distribution and, thus, needs to be solved numerically. To account for the micro evidence, we add two new ingredients to the standard calibration featuring the essence of our analysis: first, heterogeneous exposure to monetary policy shocks such that high MPC households tend to be more exposed to these shocks, and second, cognitive discounting of households with respect to aggregate shocks.

Calibration. Table 4.1 summarizes our baseline calibration. We set the discount factor β to match a steady state real rate of 2% (annualized). In contrast to Section 4.3, we now abstract from differences in time discounting, $\beta(e_{i,t}) = \beta$ for all $e_{i,t}$, such that borrowing constraints only bind for endogenous reasons. To calibrate the idiosyncratic skill process, we set $z(e_{i,t}) = e_{i,t}$ and we follow McKay et al. (2016) in assuming that $e_{i,t}$ follows an AR(1) process with autocorrelation $\rho_e = 0.966$ and variance $\sigma_e^2 = 0.033$ to match the volatility of the distribution of five-year earnings growth rates found in Guvenen et al. (2014). We then discretize this process into a three-states Markov chain using the Rouwenhorst (1995) method. 25% of households are in the lowest and the highest state, respectively, and 50% in the middle state. We set the amount of government debt to match the aggregate MPC of 0.16 out of an income windfall of 500$, as in Kaplan et al. (2018). This results in a government debt-to-annual-GDP level of 69%. We use standard parameters for our supply side. We set the the price markup to 1.2 and the Calvo probability to reset the price to 0.15 as in Christiano et al. (2011).

Table 4.1: Baseline Calibration Of the Behavioral HANK Model

Parameter	Description	Value
R	Steady State Real Rate (annualized)	2%
γ	Risk aversion	2
φ	Inverse of Frisch elasticity	2
μ	Markup	1.2
θ	Calvo Price Stickiness	0.15
ρ_e	Autocorrelation of idiosyncratic risk	0.966
σ_e^2	Variance of idiosyncratic risk	0.033
$\tau(e)$	Tax shares	[0, 0, 1]
$d(e)$	Dividend shares	$\left[\frac{0.06}{0.25}, \frac{0.18}{0.5}, \frac{0.76}{0.25}\right]$
$\frac{B^G}{4Y}$	Government debt	0.69
$\bar{m}$	Cognitive discounting	0.85

To capture fact (ii)—that higher MPC households tend to be on average more exposed to aggregate income changes induced by monetary policy—we target the estimates in Patterson (2023). Patterson (2023) finds that regressing the income elasticity of households with respect to aggregate changes in output on households' MPC yields a regression coefficient of 1.33. We match this estimate by calibrating the dividend shares the households receive. To do so, we assume that the aggregate income fluctuations are due to monetary policy shocks. We obtain a calibration that implies that households with a higher productivity receive a larger share of the dividends than households with a lower productivity. About 75% of the dividends goes to the highest productivity households. These numbers are consistent with the empirical findings in Kuhn et al. (2020).[51] To show the robustness of our results, we also consider a case in which high MPC households are even more exposed to the business cycle in Section 4.4.5.

To capture the underreaction of households to aggregate news, we set the cognitive discounting $\bar{m} = 0.85$, which corresponds to the upper bound of our empirical findings presented in Section 4.2. Since the heterogeneity in cognitive discounting seems to be small in the data (see Section 4.2), we assume that all households have the same degree of rationality in our baseline calibration. Yet, we also consider the case of heterogeneous degrees of rationality in Section 4.4.3. In Section 4.4.4, we discuss how bounded rationality with respect to *idiosyncratic* shocks affect our results.

The rest of the calibration, i.e., $\gamma = 2$, $\varphi = 2$, and the tax shares, is as in McKay et al. (2016).

[51] As MPCs are highly negatively correlated with productivity, the intuition why this leads to a higher exposure of high MPC households is exactly the same as in Section 4.3.

4.4.1 Monetary Policy

We now consider two monetary policy experiments. First, a one-time conventional expansionary monetary policy shock and second, a forward guidance shock that is announced today to take place k periods in the future. In particular, we assume that the monetary authority announces in period 0 to decrease the real interest rate by 10 basis points in period k and keeps it at its steady state value in all other periods. We follow Farhi and Werning (2019) and McKay et al. (2016) and assume that the government debt level remains constant, $B_t^G = \bar{B}^G$.

Figure 4.1: Monetary Policy and Forward Guidance

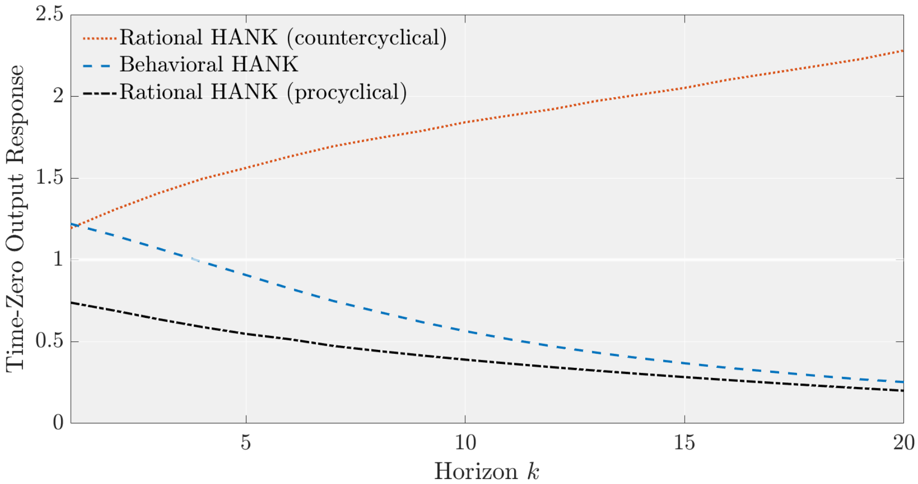

Note: This figure shows the response of total output in period 0 to anticipated one-time monetary policy shocks occurring at different horizons k, relative to the response in the representative agent model under rational expectations (normalized to 1). The blue-dashed line shows the results for the behavioral HANK model, the orange-dotted line for the rational HANK model with countercyclical inequality and the black-dashed-dotted line for the rational HANK model with procyclical inequality.

Figure 4.1 shows on the vertical axis the response of output in period 0, dY_0, to an announced real rate change implemented in period k (horizontal axis). The white horizontal line represents the response in the rational RANK model (normalized to 1). The constant response in RANK is a consequence of the assumption that forward guidance is implemented through changes in the real rate.

The blue-dashed line shows the results for the behavioral HANK model. We see that contemporaneous monetary policy has stronger effects than in RANK and the amplification is roughly 20%.[52] The intuition is the same as in the tractable model: as households with higher MPCs tend to be more exposed to aggregate income changes, monetary policy is amplified through indirect general equilibrium effects. Turning again to an AR(1)-process with a persistence of 0.6, we find that indirect effects account for 61% of the total effect in the quantitative behavioral HANK and, thus, for a large part of the transmission consistent with the findings in Kaplan et al. (2018). At the same time, the behavioral HANK model

[52] Patterson (2023) also estimates that the unequal exposure of households leads to a 20% amplification compared to an equal exposure benchmark.

does not suffer from the forward guidance puzzle, as shown by the decline in the blue-dashed line. Interest rate changes announced to take place in the future have relatively weaker effects on contemporaneous output and the effects decrease with the horizon.[53]

In contrast, the orange-dotted and the black-dashed-dotted lines highlight the tension in rational HANK models. When households with high MPCs tend to be more exposed to aggregate income fluctuations—which corresponds to $\chi > 1$ in the tractable model and which we refer to as the *countercyclical* HANK model—contemporaneous monetary policy is as strong as in the behavioral model. But with rational expectations the amplification through indirect effects extends intertemporally and results in an aggravation of the forward guidance puzzle. Indeed, we see from the orange-dotted line that the farther away the announced interest rate change takes place, the stronger the response of output today.

When, in contrast to the data, households with higher MPCs tend to be less exposed to aggregate income fluctuations—$\chi < 1$ in the tractable model and which we refer to as the *procyclical* HANK model—the rational HANK model resolves the forward guidance puzzle (see McKay et al. (2016)). But the procyclical HANK model is unable to generate amplification of contemporaneous monetary policy (see black-dashed-dotted line). Turning to an AR(1)-process, this model implies that indirect effects account only for 12% of the monetary transmission. In addition, this model has quite different policy implications, as we will see in Section 4.5.

4.4.2 Sticky Wages

As in the tractable model, we now show that our results hold when we assume that prices are fully flexible but wages are sticky. We again abstract from monopolistic competition of firms, so that prices are set to marginal costs.

Labor hours $N_{i,t}$ are determined by union labor demand. Each worker provides $N_{i,k,t}$ hours of work to a continuum of unions indexed by k. Each union aggregates efficient units of work into a union-specific task. A competitive labor packer then packages these tasks into aggregate employment services according to a CES technology and sells these services to final goods firms at price W_t. We assume that there are quadratic utility costs of adjusting the nominal wage W_{kt}. A union sets a common wage $W_{k,t}$ per efficient unit for each of its members. In doing so, the union trades-off the marginal disutility of working given average hours against the marginal utility of consumption given average consumption as in Wolf (2021). The union then calls upon its members to supply hours according to a specific allocation rule: in stationary equilibrium all households supply the same amount of hours. Outside the stationary equilibrium, we follow Auclert and Rognlie (2020) and assume the allocation rule

$$N_{i,t} = Y_t \frac{(e_{i,t})^{\zeta \log \frac{Y_t}{Y}}}{\mathbb{E}[e^{1+\zeta \log \frac{Y_t}{Y}}]}.$$

[53]We find that for our baseline calibration the behavioral HANK model resolves the forward guidance puzzle as long as $\bar{m} < 0.93$.

If $\zeta = 0$, all households supply the same amount of labor in each period. Assuming $\zeta < 0$, however, implies that the labor supply of less productive households responds more sensitively to changes in aggregate output Y_t and thus, implies countercyclical income risk. We set $\zeta = -1.2$. We further match the MPCs of 0.16 by setting the debt-to-annual-GDP level to 65%. We discretize the $e_{i,t}$ process into 11 states and as in the sticky-price model impose that only the above-median-income households pay taxes.

All in all, our setup leads to a wage Philips curve given by:

$$\pi_t^W = \kappa\left(v^{'}(N_t) - (\epsilon_n - 1)/\epsilon_n(1-\tau_t)\frac{W_t}{P_t}u^{'}(C_t)\right) + \beta\pi_{t+1}^W, \tag{4.25}$$

where $\epsilon_n = 11$ is the elasticity of substitution between differentiated labor supply and $\kappa = 0.1$ is the slope of the wage Philips Curve.

Figure 4.2 shows the effects of conventional monetary policy shocks and of forward guidance shocks on output at time 0 in our sticky-wage behavioral HANK model (blue-dashed line) as well as for the rational HANK model with sticky wages (orange-dotted line). We see that our results are robust. Whereas in both models, contemporaneous monetary policy is to a large share transmitted through indirect effects, the behavioral HANK model rules out the forward guidance puzzle whereas it is aggravated in the rational HANK model compared to the representative agent model.

Figure 4.2: Sticky Wages

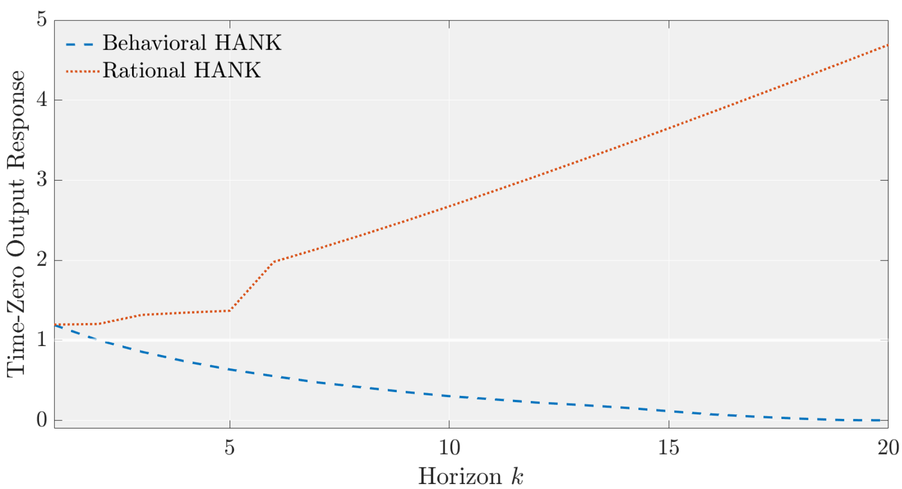

Note: This figure shows the response of total output in period 0 to anticipated i.i.d. monetary policy shocks occurring at different horizons k for the case in which prices are flexible and wages are sticky.

4.4.3 Heterogeneous Cognitive Discounting

So far, we have assumed that all households exhibit the same degree of rationality. Yet, as we showed in Section 4.2, while underreaction is found across all income groups, the data suggests that higher income households deviate somewhat less from rational expectations.

To model this, we assume that a household's rationality is a function of her productivity level e: $\bar{m}(e = e_L) = 0.8$, $\bar{m}(e = e_M) = 0.85$ and $\bar{m}(e = e_H) = 0.9$.

This parameterization serves three purposes: first, the lowest-productivity households exhibit the largest deviation from rational expectations and the degree of rationality increases monotonically with productivity. Second, the average degree of bounded rationality remains 0.85 such that we can isolate the effect of heterogeneity in bounded rationality from its overall level. And third, this is a rather conservative parameterization—both in terms of the degree of heterogeneity and in the level of rationality—compared to the results in the data which points towards lower levels of rationality across all households and less dispersion. We discuss an alternative calibration—one in which a subgroup of households is fully rational—in Appendix D.5.

Figure 4.3: Heterogeneous $\bar{m}$ and Monetary Policy

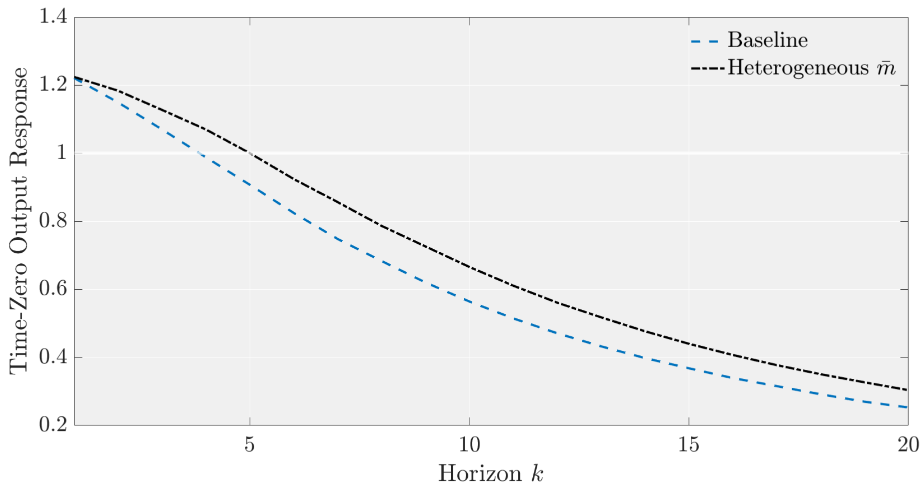

Note: This figure shows the response of total output in period 0 to anticipated i.i.d. monetary policy shocks occurring at different horizons k for the baseline calibration with $\bar{m} = 0.85$ for all households (blue-dashed line) and for the model in which households differ in their levels of cognitive discounting (black-dashed-dotted line).

Figure 4.3 compares the model with heterogeneous degrees of bounded rationality (black-dashed-dotted line) to our baseline quantitative behavioral HANK model (blue-dashed line) for the same monetary policy experiments as above. The effect of a contemporaneous monetary policy shock is practically identical across the two scenarios consistent with the insight that amplification of contemporaneous monetary policy is barely affected by the degree of rationality. At longer horizons, however, monetary policy is more effective in the economy in which households differ in their degrees of rationality.

There are two competing effects: first, high productivity households are now more rational such that they react stronger to announced future changes in the interest rate compared to the baseline which increases the effectiveness of forward guidance. Second, low productivity households are less rational which tends to dampen the effectiveness of forward guidance. Yet, a large share of low productivity households are at their borrowing constraint and, thus, they do not directly react to future changes in the interest rate anyway while most of the high productivity households are unconstrained. Hence, the first effect dominates and forward guidance is more effective compared to the baseline model. Overall, however, the differences across the two calibrations are rather small. As we show in Appendix D.5, even when the highest productivity households are fully rational the forward guidance puzzle is resolved and the effects of forward guidance vanish quite quickly with the horizon.

4.4.4 Non-Rational Expectations about Idiosyncratic Shocks

Up to now, we have assumed that households are fully rational with respect to aggregate shocks but that households are perfectly rational with respect to their own idiosyncratic risk. Yet, recent evidence suggests that professional forecasters (Kučinskas and Peters (2022)) and firms (Born et al. (2022)) show patterns of *overreaction* with respect to individual shocks. We now show how simultaneous underreaction to aggregate shocks and overreaction to idiosyncratic shocks affect our results.

To do so, we extend our model and now assume that households overpredict the persistence of their idiosyncratic risk, that is $\tilde{\rho}_e > \rho_e$, where ρ_e denotes the persistence of their actual risk and $\tilde{\rho}_e$ denotes the perceived persistence. In particular, we consider $\tilde{\rho}_e = 0.976$ (instead of $\rho_e = 0.966$). We verify that this implies overreaction to individual news by running the Coibion and Gorodnichenko (2015) regressions (see equation (4.7)) on model-simulated data with the idiosyncratic productivity being the forecasted variable, $e_t = x_t$. This yields a regression coefficient that is negative, indicating overreaction (Bordalo et al. (2020)).[54]

The orange-dashed line in Figure 4.4 shows the effects of monetary policy and forward guidance shocks for our extended model where households also misperceive their individual risk. The main take-away is that our results are very robust to this extension: monetary policy is amplified through indirect, general-equilibrium effects and the effectiveness of forward guidance decreases with the horizons. Quantitatively, the results are pretty similar to our baseline model (blue-dashed line), in which households are fully rational with respect to their idiosyncratic risk. Qualitatively, the effectiveness of forward guidance becomes weaker in the model with overreaction to idiosyncratic shocks. To understand this, recall the aggregate IS equation in our tractable model, equation (4.18), but replace the actual idiosyncratic risk $1 - s$ with the perceived risk $1 - \tilde{s}$:

$$\widehat{y}_t = \tilde{\psi}_f \mathbb{E}_t \widehat{y}_{t+1} - \psi_c \frac{1}{\gamma} \widehat{r}_t,$$

with $\tilde{\psi}_f \equiv \bar{m}\left[1 + (\chi - 1)\frac{1-\tilde{s}}{1-\lambda\chi}\right]$. Overreaction is captured by $1 - \tilde{s} < 1 - s$. We can thus directly see that $\tilde{\psi}_f < \psi_f$ when $\chi > 1$, which dampens the effects of forward guidance.

[54]We simulate the model for 1000 households over 500 periods and do this 100 times. For each simulation, we estimate b^{CG} using one-quarter ahead forecast errors. The mean estimate is -0.01 in the case with $\tilde{\rho}_e = 0.976$. Thus, overreaction is quite small, consistent with the empirical findings in Afrouzi et al. (2022) for highly persistent AR(1) processes.

Intuitively, even though unconstrained households know that they will benefit more from an expansionary forward guidance shock in case they become hand-to-mouth, they under-predict the probability of becoming hand-to-mouth, and hence, forward guidance is further dampened. The same mechanism is at work in our full model, which explains the results in Figure 4.4.

Figure 4.4: Non-Rational Expectations about Idiosyncratic Shocks

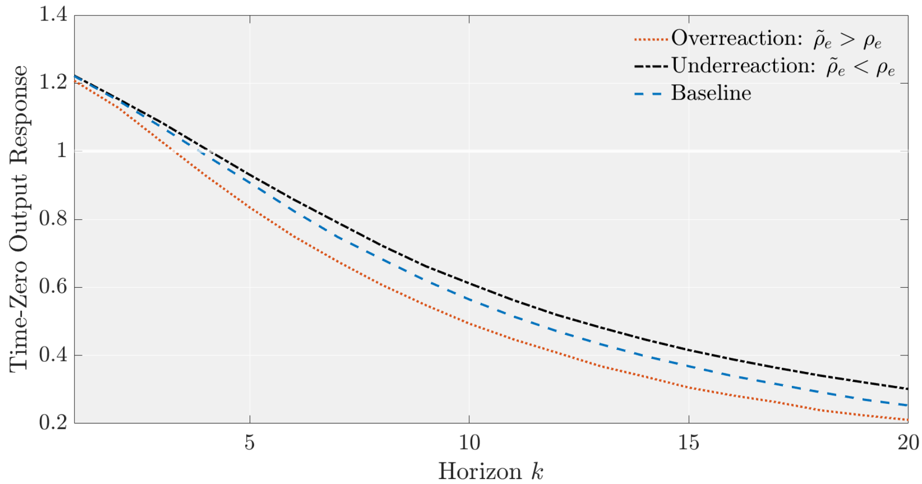

Note: This figure shows the response of total output in period 0 to anticipated one-time monetary policy shocks occurring at different horizons k, relative to the response in the representative agent model under rational expectations (normalized to 1). The orange-dotted line shows the overreaction case with $\tilde{\rho}_e = 0.976 > \rho_e$, the black-dashed-dotted line the underreaction case with $\tilde{\rho}_e = 0.95 < \rho_e$, and the blue-dashed line our baseline model in which households correctly perceive the persistence of their idiosyncratic risk, $\tilde{\rho}_e = \rho_e = 0.966$.

For completeness, we also consider a case in which households underreact to idiosyncratic shocks. In particular, we set $\tilde{\rho}_e = 0.95$ and, thus, $\tilde{\rho}_e < \rho_e$. The black-dashed-dotted line in Figure 4.4 shows that also in this case, our main results remain very robust and the quantitative differences to our baseline model are small. Qualitatively, the effectiveness of forward guidance is now a bit less dampened consistent with our intuition above: as households underestimate the persistence of their idiosyncratic productivity, they are more eager to precautionary save and, thus, they react stronger to the relaxation in their precautionary savings risk induced by news about future interest rate decreases. We also run a robustness check with an even more extreme degree of underreaction towards individual news ($\tilde{\rho}_e = 0.85$, not shown) and we find that even in this case, the forward guidance puzzle remains solved.

The take-aways from this section are that our results are robust to allowing for households to deviate from rational expectations about their idiosyncratic risk and that in the probably empirically-relevant case of overreaction, the effects of forward guidance are even further dampened.

4.4.5 Further Results

Stability at the effective lower bound. To test the stability of the model at the effective lower bound—fact (iv)—we consider a shock to the discount factor that pushes the economy to the ELB for 8 periods, in the behavioral and the rational model. After that the shock jumps back to its steady state value. Consistent with the tractable model, the recession in the rational model is substantially more severe. While output drops on impact by 5% in the behavioral model, it drops by 10% in the rational model (see Appendix D.5.2 for details).

Unequal exposure: more extreme calibration. In our baseline calibration, we target the finding from Patterson (2023) that a linear regression of households' income elasticity to GDP on their MPC yields a coefficient of 1.33. In Appendix D.5.1, we show that our results remain robust when we target a more extreme coefficient of 2. In this case, the initial amplification through indirect effects becomes stronger, but the model still resolves the forward guidance puzzle and, thus, is able to account for fact (i) - (iv) simultaneously.

4.5 Policy Implications of Inflationary Supply Shocks

Having established that the behavioral HANK model is consistent with recent facts about the transmission and effectiveness of monetary policy, we now use the model to revisit the policy implications of inflationary supply shocks. We uncover a novel amplification channel of these shocks that is absent in existing models as it arises due to the interaction of the unequal exposure of households to monetary policy and the behavioral friction, and thus, exactly through the model ingredients that allow the model to simultaneously account for facts (i)-(iv).

Many advanced economies have recently experienced a dramatic surge in inflation and at least part of this is attributed to disruptions in production, such as supply-chain "bottlenecks" (see, e.g., di Giovanni et al. (2022)). We model these disruptions as a negative total factor productivity (TFP) shock. Production of intermediate-goods firm j is now given by $Y_t(j) = A_t N_t(j)$, where A_t is total factor productivity following an AR(1)-process, $A_t = (1-\rho_A)\bar{A} + \rho_A A_{t-1} + \varepsilon_t^A$, and ε_t^A is a zero-mean i.i.d. shock, $\bar{A}$ the steady-state level of TFP and ρ_A the persistence of A_t which we set to $\rho_A = 0.9$. Each firm can adjust its price with probability 0.15 in a given quarter and we assume that firms have rational expectations to fully focus on the role of bounded rationality on the household side (we discuss the case with behavioral firms later).

Government debt is time-varying and total tax payments, T_t, follow the debt feedback rule, $T_t - \bar{T} = \vartheta \frac{B_{t+1}^G - \bar{B}^G}{\bar{Y}}$, where we set $\vartheta = 0.05$. We start with the case in which monetary policy follows a simple Taylor rule (4.10) with an inflation coefficient of 1.5. Later on, we discuss the case in which monetary policy follows a strict inflation-targeting rule and implements a zero inflation rate in all periods.

The size of the shock is such that output in the model with fully-flexible prices, complete markets and rational expectations—what we from now on call *potential output*—decreases by 1% in terms of deviations from its steady state. We normalize the leisure parameter in the complete markets, flexible price model such that it has the same steady state output as our behavioral HANK model. The *output gap* is then defined as the difference between actual output and potential output divided by steady state output.

Figure 4.5 shows the impulse-response functions of output, the output gap, inflation, nominal interest rates, government debt and the consumption Gini index as a measure of inequality after the negative supply shock. The blue-dashed lines show the responses in the behavioral HANK model, the orange-dotted lines in the rational HANK model, and the black-solid lines in RANK. We assume government debt to be constant in RANK.[55]

Figure 4.5: Inflationary supply shock: Taylor rule

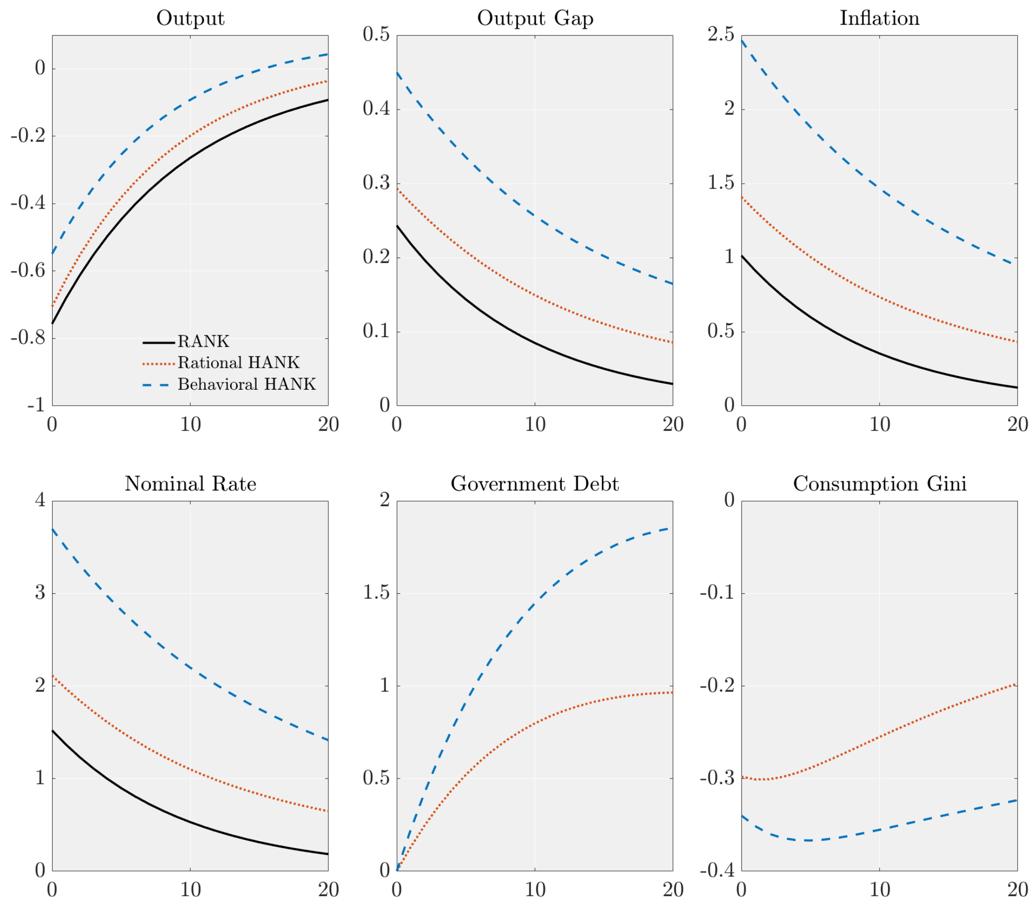

Note: This figure shows the impulse responses after a productivity shock for the case where monetary policy follows a Taylor rule. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per annual-GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

[55] As we implement taxes such that they do not show up in the first-order conditions of household, in the RANK version of our model Ricardian equivalence holds and, thus, the path of debt does not matter anyway.

Qualitatively, the impulse responses are the same across all models: in response to the supply shock, monetary policy increases the nominal interest rate, which pushes down output. Yet output falls by less than potential output, leading to positive output gaps which pushes up inflation. Yet, *quantitatively* there are large differences across the models. In particular, the increase in inflation is roughly 2.5 times as large in the behavioral HANK model compared to RANK and 1.7 times as strong as in the rational HANK model even though the (nominal and real) interest rate increases most strongly in the behavioral HANK model. The reason is a novel amplification channel due to household heterogeneity, cognitive discounting and the interaction of the two. The positive output gap increases wages and decreases profits relative to the outcome without nominal rigidities in the same way as expansionary policy shocks in Sections 4.3 and 4.4 do. This redistributes on average towards lower income and higher MPC households which further increases the output gap and inflation. In addition, the higher expected real interest rates in response to the negative supply shock lead to a negative deviation of expected consumption from its stationary equilibrium counterpart. In the behavioral HANK model, households cognitively discount the expected higher interest rates and, hence, their consumption expectations decrease by less. As a result, households decrease today's consumption by less compared to fully rational households. This further increases the output gap which amplifies the redistribution to high MPCs households which again amplifies the increase in the output gap until the economy ends up in an equilibrium with a higher output gap and higher inflation.

While inflation and the output gap increase substantially, consumption inequality decreases both in the rational as well as in the behavioral HANK model and it decreases even more in the behavioral model (see lower-right panel in Figure 4.5). While higher interest rates redistribute to relative consumption-rich households, this effect on consumption inequality is dominated by the increase in the output gap which redistributes to relatively consumption-poor households. Finally, the higher real interest rates increase the cost of government debt which is (partly) financed by issuing more debt. Thus, the government debt level increases, especially in the behavioral HANK model where the increase in real interest rates is larger.

Given the larger sensitivity of inflation to supply shocks due to this novel amplification channel, our model may hence offer a (partial) explanation for why many advanced economies have seen large inflation increases following the Covid-19 pandemic. In particular, our model predicts that when the shock redistributes towards high-MPC households and when the central bank's response to inflation is underpredicted, inflationary supply shocks can lead to a substantial increase in inflation.

Decomposition of the amplification channel. How much of the additional inflation increase in the behavioral HANK model compared to RANK is due to the underlying heterogeneity, how much is due to cognitive discounting and how much is due to the interaction of the two? Figure 4.6 decomposes the amplification channel into these three components. It shows the additional inflation increase in the behavioral HANK model compared to the inflation increase in the RANK benchmark and its components. The black-solid line shows the inflation response in RANK. The orange-shaded area denotes the additional inflation increase that arises solely due to household heterogeneity. The blue-shaded area the fraction of the overall increase due to cognitive discounting alone. Thus, the gray-shaded area captures the additional inflation increase that is due to the interaction of household heterogeneity and cognitive discounting.

Under our baseline calibration, this complementarity amounts to about 27% on impact of the inflation response in RANK (the inflation response in RANK is 1 percentage point on impact). As the additional increase in the behavioral HANK is about 1.45 percentage points, the complementarity explains about 19% of the *additional increase*. Cognitive discounting alone accounts for 27% while household heterogeneity accounts for 54% of the amplification over RANK. Figure D.9 in Appendix D.6.1 considers an alternative calibration of the discounting parameter where we set it to 0.6 and thus the lower bound of the empirical estimates instead of the upper bound. In this case, inflation increases more than 3.5 times as much as in the RANK model with the interaction between household heterogeneity and cognitive discounting accounting for 1/3 of the initial amplification and the interaction itself accounts for a larger share of the overall additional increase than the underlying heterogeneity.

Figure 4.6: Decomposition of the Additional Inflation Increase

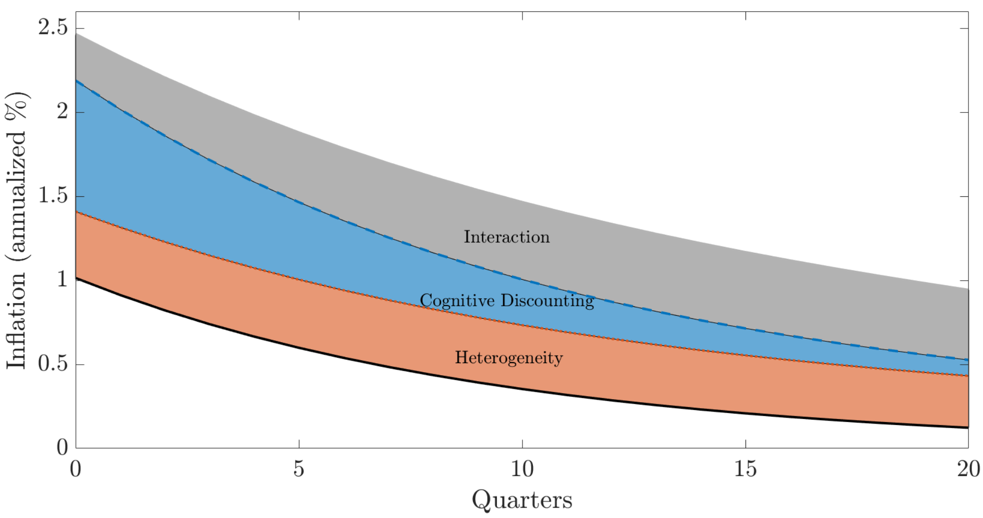

Note: This figure shows the decomposition of the additional inflation increase in the behavioral HANK model compared to the rational RANK model. The orange-shaded area represents the additional increase that is solely due to the heterogeneous exposure of households, the blue area the increase due to cognitive discounting and the gray area the additional increase that is due to the interaction of heterogeneity and cognitive discounting.

Supply vs. demand shocks. The fact that the behavioral HANK model amplifies persistent supply shocks more than the rational HANK model is in contrast to persistent demand shocks. While both the underlying heterogeneity and bounded rationality amplify persistent supply shocks, these both model features work in opposite direction in response to persistent demand shock. For example, in response to an expansionary monetary policy shock, the heterogeneous exposure of households amplifies the effects of the shock. The high-MPC households benefit more strongly from the shock which triggers an amplification of the shock, as discussed extensively in Sections 4.3 and 4.4. Cognitive discounting, however, would dampen the effect because households would discount the persistent decrease in the interest rate. As a result, persistent demand shocks are less strong in the behavioral HANK model compared to the rational HANK model: an expansionary monetary shock of 1 percentage point with a persistence of 0.6 increases inflation on impact by by 1.24pp. (annualized) in the behavioral HANK while it does by 1.44pp. in the rational HANK model.

Strict inflation targeting. What if monetary policy reacts more hawkish to inflation? Figure 4.7 shows the limiting case, in which monetary policy follows a strict inflation targeting rule and, hence, keeps inflation at zero at all times. We see that the output responses are almost indistinguishable across the two models and practically identical to the fall in potential output such that the output gap is essentially zero.

Yet, the reaction of monetary policy differs significantly across the two models. The nominal (and real) interest rate in the behavioral HANK model increases twice as much on impact as in the rational HANK model. The reason is that behavioral households cognitively discount the future higher interest rates that they expect due to the persistence of the shock. Hence, these expected higher future rates are less effective in stabilizing inflation today. Thus, to induce zero inflation in every period, monetary policy needs to increase interest rates by more than in the rational HANK model, in which the expected future interest rate hikes are very powerful. As this line of reasoning applies in each period, the interest rate in the behavioral HANK model remains above the interest rate in the rational model.

Raising interest rates increases the cost of debt for the government which it finances in the short run by issuing additional debt. The bottom-middle panel in Figure 4.7 shows that government debt in the behavioral model increases by more than twice as much as in the rational model and by more than when monetary policy follows a simple Taylor rule. Thus, the fiscal footprint of monetary policy is larger because monetary policy needs to respond more strongly to counteract the inflationary pressures in the behavioral model.

On top of the stronger increase in government debt and interest rates, consumption inequality increases more strongly in the behavioral model compared to the rational model. The reason is that along the wealth distribution, increases in the real interest rate redistribute to wealthier households and, hence, to households who already tend to have a higher consumption level. As the increases in the real interest rate are higher in the behavioral HANK model, these redistribution effects are more pronounced. Because monetary policy fully stabilizes inflation and the output gap, dividends and wages fall by the same relative amount after the productivity shock, such that each household's labor and dividend income falls by the same amount. Hence, the redistribution channels present in Sections 4.3 and

Figure 4.7: Inflationary supply shock: strict inflation-targeting

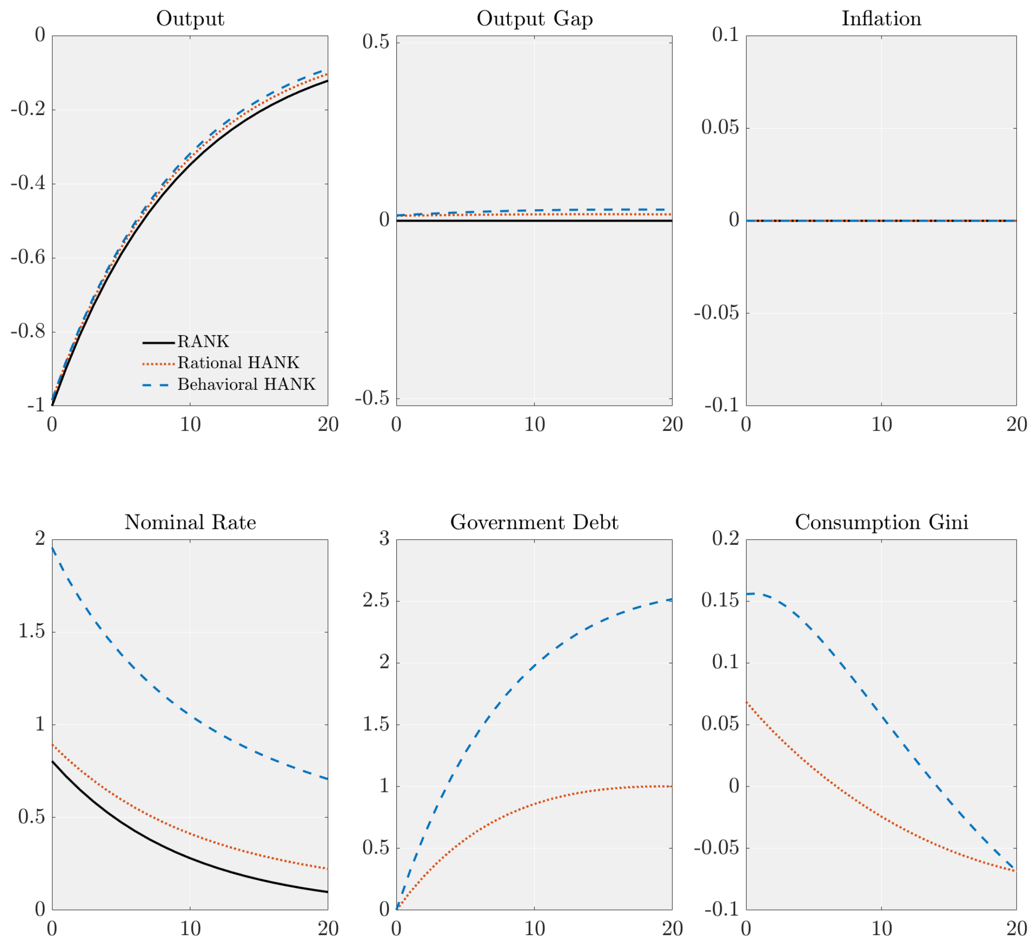

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

4.4 after policy shocks are muted here. Put differently, monetary policy turns off the amplification mechanism that works through the unequal exposure of households when implementing zero inflation.

Overall, our model suggests that accounting for facts (i)-(iv) simultaneously has important implications for policy. In particular, there is a strong trade off for monetary policy following an inflationary supply shock. Simply following a Taylor rule and thus, not responding very aggressively to the inflationary pressures, can lead to a significant increase in inflation through the mutual reinforcement of households' unequal exposure to the overheating of the economy as well as households' cognitive discounting of the monetary authority's future response to inflation. Counteracting these forces and implementing a zero inflation rate, however, requires a much stronger monetary policy response which, in turn, leads to a strong increase in government debt and inequality.

Comparison to the procyclical HANK model. One of the reasons why the behavioral HANK model amplifies supply shocks is that it is less responsive to expected future interest rates. A natural question is then: how do its policy implications compare to those derived in rational HANK models that are calibrated to resolve the forward guidance puzzle? As shown in Section 4.4, when all households receive an equal share of the dividends, the rational model can resolve the forward guidance puzzle (McKay et al. (2016)). This implies that households with high MPCs benefit less from income increases induced by monetary policy, thereby violating fact (ii).

Figure D.10 in the appendix shows that this "procyclical" rational HANK model predicts a much weaker response of inflation to the same supply shock in the case of a standard Taylor rule. The reason is that now the positive output gap redistributes on average to high-income and low MPC households which further dampens aggregate demand. In other words, this model features a dampening channel compared to RANK after supply shocks instead of an amplification channel as in the behavioral HANK model.[56]

The two models also differ in terms of their cross-sectional implications: in the procyclical HANK model, consumption inequality increases strongly whereas it decreases in the behavioral HANK model.

Behavioral firms. In Appendix D.6.3, we discuss the case in which firms cognitively discount the future in the same way as households. The increase in inflation when monetary policy follows a Taylor rule is somewhat muted whereas the increase in the output gap is amplified compared to the case in which firms are rational. The reason is that firms discount the increase in their future marginal costs and thus increase their prices not as strongly. According to the Taylor rule this then leads to a smaller increase in interest rates so that households consume more, leading to an increase in demand and thus, the output gap.

Cost-push shocks. So far, we have focused on the inflationary pressure coming from negative TFP shocks. We show in Appendix D.6.4 that if the inflationary pressure comes from a cost-push shock instead, the monetary and fiscal implications are very similar. Inflation and the output gap increases much more in the behavioral HANK model compared to the rational HANK model although interest rates increases more. Accordingly, if the central bank wants to fully stabilize inflation, it needs to raise interest rates much more strongly in the behavioral HANK model than in the rational HANK model to fully stabilize inflation. This pushes up the government debt level, especially in the behavioral HANK model.

[56]Another take-away is that for a given persistent demand shock, the behavioral HANK model and a recalibrated version of the procyclical HANK model could be observationally equivalent in terms of the output and inflation response. Yet, these two models then differ drastically after supply shock.

4.6 Conclusion

In this paper, we develop a new framework for business-cycle and policy analysis: the behavioral HANK model. To arrive at our framework, we introduce bounded rationality in the form of cognitive discounting and household heterogeneity into a New Keynesian model. The model can account for recent empirical findings on the transmission mechanisms of monetary policy. In particular, households with higher marginal propensities to consume tend to be more exposed to changes in aggregate income that are induced by monetary policy, leading to an amplification of conventional monetary policy through indirect effects. Simultaneously, the model rules out the forward guidance puzzle and remains stable at the effective lower bound. The model thus overcomes a tension in existing models with household heterogeneity: when accounting for the underlying heterogeneity, these models tend to aggravate the forward guidance puzzle and the instability issues at the lower bound. Both, bounded rationality and household heterogeneity, are crucial to arrive at our results.

Simultaneously accounting for these facts matters greatly for the model's policy implications. In particular, we uncover a new amplification mechanism of inflationary supply shocks through cognitive discounting and the unequal exposure of households. After a negative productivity shock the behavioral HANK model predicts a substantially larger inflation increase. If the monetary authority wants to stabilize inflation after such an inflationary supply shock, it needs to hike the nominal interest rate much more strongly than under rational expectations which leads to a strong increase in government debt and inequality.

Given its consistency with empirical facts about the transmission of monetary policy, the behavioral HANK model provides a natural laboratory for both business-cycle and policy analysis. Our framework can also easily be extended along many dimensions, some of which we have explored in the paper, whereas others are left for future work.

Chapter 5

Bad luck or bad decisions? Macroeconomic implications of persistent heterogeneity in cognitive skills and overconfidence

with Oliver Pfäuti and Jonathan Zinman

Abstract

Business cycle models often abstract from persistent household heterogeneity, despite its potentially significant implications for macroeconomic fluctuations and policy. We show empirically that consumers' likelihood of being persistently financially constrained decreases with cognitive skills and increases with overconfidence thereon. Guided by this and other micro evidence, we add persistent heterogeneity in cognitive skills and overconfidence to an otherwise standard HANK model. Overconfidence proves to be the key innovation, driving households to spend instead of precautionary save and producing empirically realistic wealth distributions and hand-to-mouth shares and MPCs across the income distribution. We highlight implications for various fiscal policies.

Keywords: Household Heterogeneity; Cognitive Skills; Overconfidence; Financial Constraints; Fiscal Policy; HANK
JEL-Classification: D91; E21; E62; E71; G51.

5.1 Introduction

Heterogeneity in households' savings behavior and financial situations has significant implications for macroeconomic fluctuations and policy design.[57] Yet it remains standard practice in macro modeling to assume ex-ante identical households and account for heterogeneity only in shock realizations: Households are wealthy or poor only because of good luck or bad luck, abstracting from choices linked to fundamental and persistent dimensions of heterogeneity across households.[58]

One important dimension of fundamental heterogeneity is cognitive skills, which has been linked empirically to: differences in economic growth across space and time (Hanushek and Woessmann (2008)), households' inflation expectations (D'Acunto et al. (2019a), D'Acunto et al. (2019b)), responses to changes in incentives (D'Acunto et al. (2023)), financial mistakes (Agarwal and Mazumder (2013)), and strong negative relationships between behavioral biases and income (e.g., Stango and Zinman (2020), Chapman et al. (2023)). Links between cognitive skills and savings behavior are less well understood, particularly for the sorts of behaviors and outcomes featured in macro modeling (e.g., hand-to-mouth a.k.a. HtM status). And the macroeconomic implications of any such link between cognitive skills heterogeneity and HtM heterogeneity are largely unexplored.

We start by using micro data, from a nationally representative sample of working-age U.S. consumers, to develop several new facts about how cognitive skills, beliefs, and household financial situations are related. We find that the likelihood of being persistently HtM, measured in various ways, decreases sharply with cognitive skills. But allowing for cognitive skills heterogeneity alone is unlikely to help macro models fit the data, as we later formalize, because permanently low-productivity households will still tend to save their way out of HtM status if they are classically rational. This motivates considering beliefs as well, starting with how consumers perceive their own cognitive skills.

We show that persistent overestimation of one's own skills is prevalent (as in the high-stakes setting of Huffman et al. (2022)) and differs across consumers: overconfidence correlates strongly and negatively with cognitive skills. Overconfident consumers are also about 1.2 times as likely than their well-calibrated counterparts to be persistently overly-optimistic about their future financial situations (measured using standard consumer sentiment forecasts and realizations), suggesting that lower-skilled consumers may be HtM at least in part because of their overconfidence. Consistent with this conjecture, we find strong correlations between persistent overconfidence and our measures of persistent HtM status.

Guided by these micro findings, we add persistent heterogeneity in cognitive skills and

[57] See, e.g., Werning (2015), Kaplan et al. (2018), Auclert (2019), Bayer et al. (2020a), Luetticke (2021), Hagedorn et al. (2019), Patterson (2023), Almgren et al. (2022), Holm et al. (2021) on shock transmission and policy efficacy, and Dávila and Schaab (2022), McKay and Wolf (2022), Bhandari et al. (2021), Bilbiie (2021), Smirnov (2022), Acharya et al. (2023), Yang (2022) on optimal policy design.

[58] Important exceptions include models allowing for heterogeneity in preferences (e.g., Auclert et al. (2020), Aguiar et al. (ming), Kaplan and Violante (2022)), as we discuss below.

perceptions thereof to an otherwise standard heterogeneous agent New Keynesian (HANK) model with incomplete markets, idiosyncratic productivity risk, borrowing constraints, and a nominal rigidity in the form of sticky wages. This framework allows us to unpack potential mechanisms underlying our empirical findings, and to derive macroeconomic implications of heterogeneity in cognitive skills and overconfidence. We model cognitive skills heterogeneity as differences in average labor market productivity and overconfidence as overweighting the probability of reaching a better productivity state and underweighting the probability of reaching a worse state. Motivated by micro data on the prevalence of persistent overconfidence and the strong correlation between cognitive skills and overconfidence, and in the interest of parsimony, we calibrate our baseline model such that 62% of households are high-skilled with well-calibrated beliefs about future productivity while the remaining 38% are low-skilled and overconfident. We also calibrate the parameter governing the degree of overconfidence, by matching our finding that overconfident households are about 1.2 times as likely to be overly-optimistic about their future financial situations than rational households.

Accounting for heterogeneity in cognitive skills and overconfidence substantially improves the model's empirical fit. In contrast to standard one-asset HANK models and to a HANK model with heterogeneity in skills but not in beliefs about them, our model jointly matches total wealth in the economy, high HtM prevalence, and an average quarterly marginal propensity to consume (MPC) in the consensus range of 15-25% (e.g., Jappelli and Pistaferri (2010), Havranek and Sokolova (2020)). This holds even when all wealth is liquid and held in a single asset.[59]

Existing one-asset HA(NK) models struggle to match these data moments jointly because if the supply of assets is large enough to match the average wealth in the economy, the price of the asset is so low that almost all households accumulate a sufficient buffer stock to make the borrowing constraint nonbinding (Auclert et al. (ming), Kaplan and Violante (2022)). This makes HtM status counterfactually rare and implies that most households have low MPCs. Consequently, standard models produce an average MPC that is too low.

Our model achieves reconciliation because overconfident households underestimate their insurance needs and consequently perceive the price of the asset as too high to merit accumulating a sufficient buffer stock. Even when the supply of assets is high, many overconfident households choose to do little if any precautionary saving and often end up being HtM, consistent with our empirical findings. HtM status in our model thus is often due at least partly to "bad decisions" and not only to the "bad luck" that drives standard models. Our results are driven by differences in overconfidence rather than by differences

[59]Our model also accounts well for other untargeted wealth inequality statistics. It produces more and more empirically realistic inequality than its rational counterpart, better matching empirical wealth shares—e.g., of the top 10% or the bottom 50%. Moreover, our model does not suffer from the "missing middle" problem (Kaplan and Violante 2022) of an implied wealth distribution that is too polarized compared to the data. E.g. a standard one-asset HANK model predicts median wealth that is about an order of magnitude smaller than the data's. Our model matches this (untargeted) moment well.

in skills: removing heterogeneity in overconfidence from the model by imposing rational beliefs for all households, while retaining heterogeneous average skill levels, fails to match the average MPC and delivers very few HtM households.

A standard practice for better reconciling HANK models with the data is to introduce a second, illiquid asset that can be adjusted only infrequently (Kaplan and Violante (2014), Kaplan et al. (2018), Bayer et al. (2019), Auclert et al. (ming)). This approach produces a liquidity premium that is arguably too high, as discussed in Kaplan and Violante (2022). We show that a two-asset version of our model can fit the data with a substantially lower liquidity premium, because overconfident households underestimating their individual income risk implies that they also underestimate the shadow value of future liquidity and thereby put downward pressure on the equilibrium liquidity premium.

In contrast to standard models, our model also generates empirically realistic shares of HtM households throughout the income distribution, even though we do not explicitly target this. Because overconfidence is a key predictor of HtM status, our model produces significant shares of higher-income HtM households, in line with the data. Standard models produce either far too few HtM households throughout the income distribution (when calibrated to match average wealth) or far too much HtM polarization by income (when directly targeting the average MPC). The reason is that standard models match the average MPC by making practically all low-income households HtM—a side effect of households being HtM due only to "bad luck".

Our model thus requires only one additional parameter—and no free parameters, as we discipline overconfidence using our new survey evidence—to substantially improve the performance of existing HANK models. The mechanism that allows us to better match key features of micro and macro data—lower-skilled households' undersaving due to overconfidence about their future financial situations—generates important and distinct implications for macroeconomic policies as well.

We start our policy analysis by considering unexpected transfer payments intended to stimulate private consumption by targeting households with high MPCs. Given the difficulty of empirically identifying the general-equilibrium effects of transfer policies, they are usually evaluated using models. For untargeted transfers, these models sensibly match the average MPC (e.g., Kaplan and Violante (2014), Wolf (2021)). But with targeted transfers, average MPCs are no longer sufficient to assess the aggregate effects: it is the *distribution* of MPCs, across targeted vs. non-targeted groups, that matters. We illustrate this by modeling a stimulus payment targeted to the bottom income quartile and estimate a transfer multiplier of 0.9 in general equilibrium. This contrasts with standard HANK models, which either under- or overestimate the share of HtM households in low-income groups and thus under- or overestimate the average MPC of transfer recipients: the low-MPC standard HANK model predicts a multiplier of 0.5 and its low-wealth, high-MPC counterpart a multiplier of 2.4.

Next we show that heterogeneity in overconfidence also has implications for fiscal policies that more directly impact household self-insurance decisions in steady-state. The key mechanism is that our financially constrained households are mostly overconfident and hence value additional insurance less than the constrained and rational households in standard models.

First, we show that providing public insurance through minimum income benefits does not crowd out private precautionary savings as strongly as predicted by rational models. Overconfident households undervalue this insurance because they underestimate their probability of reaching bad income states and therefore reduce any existing buffer stock only mildly. The introduction of minimum income benefits thus only weakly increases the steady-state share of HtM households and the equilibrium real interest rate in our model, in contrast to standard models.

Second, we consider indirect insurance provision through government debt issuance (e.g., Woodford (1990), Aiyagari and McGrattan (1998)). Higher government debt levels reduce households' self-insurance cost by reducing the cost of liquid assets. But the induced increase in precautionary saving is muted in our model because overconfident households undervalue the insurance function of cheaper assets. Thus even at high public debt levels, many overconfident households do not save themselves out of being constrained, the HtM share remains high, and the wealth share of the bottom 50% remains stubbornly low. In a standard model, low-wealth households are eager to save themselves away from the borrowing constraint and increase their saving strongly in response to cheaper liquidity. This drives down the HtM share strongly and increases the wealth share held by the bottom half of the distribution. These contrasting effects have normative implications as well: the optimal government debt level is substantially lower with heterogeneous overconfidence, irrespective of whether we consider a model in which households can only save in government bonds or also in productive capital.

Related literature. We contribute to five strands of literature. One considers how cognitive skills heterogeneity affects the macroeconomy. So far, this literature is largely empirical and focused on growth (Hanushek and Woessmann (2008)). D'Acunto et al. (2019a), D'Acunto et al. (2019b), and D'Acunto et al. (2023) bring cognitive skills heterogeneity to the empirical study of economic fluctuations, showing it plays key roles in how households form their inflation expectations and respond (or not) to information and incentives provided by policy interventions. We empirically link heterogeneity in cognitive skills to heterogeneity in forecasted and realized financial situations, including HtM status. We then build a model capturing key features of that micro heterogeneity and use it to quantitatively study macro dynamics, the wealth distribution, and policy design and effectiveness. Altogether, we show that accounting for heterogeneity in beliefs about cognitive skills can greatly improve a model's ability to fit the micro and macro data.

A second strand considers potential psychological sources of liquidity or poverty traps and their macro implications. Work on aspirations as reference points (Dalton et al. 2016;

Genicot and Ray 2017) has focused on how excessive pessimism can dampen growth, while we focus on how excessive optimism affects stabilization and macroeconomic policies. Sergeyev et al. (2024) consider how financial stress and naivete about financial stress can create persistent financial constraints and impact wealth inequality and fiscal multipliers. We consider a different decision making mechanism than work on aspirations or stress, focusing on biased beliefs rather than behavioral preferences or the neglect of one or more key parameters. Our mechanism is relatively easy to validate empirically and incorporate into an otherwise standard quantitative model.

A third strand focuses on differences between perceived vs. actual idiosyncratic labor market risk. So far, this literature has focused on various beliefs about the labor market and a subset of important macro applications. Balleer et al. (2022) show that working-age individuals in the U.S. are "vastly over-optimistic about their own labor market prospects" (p. 1). Mueller et al. (2021) find optimistic bias about job-finding rates, especially for the long-term unemployed, and little evidence for downward revision of these beliefs when remaining unemployed. Wang (2023) shows how calibrating a standard incomplete-markets model to consumers' perceived rather than actual income risk is better able to account for observed wealth inequality.[60] Our contributions are uncovering the role of cognitive skills heterogeneity in shaping biased perceptions about risk and future financial situations, and building a general equilibrium model that can jointly fit key features of micro and macro data and quantitatively evaluate and guide policy.

Fourth, we contribute to the development of macro models seeking to use insights from behavioral economics to improve predictive and prescriptive power. Most work in this vein focuses on a representative behavioral agent.[61] Behavioral HANK models tend to allow for heterogeneity only in the budget constraint, with a homogeneous behavioral or information friction about an aggregate variable only.[62] Pfäuti and Seyrich (2023) study a case of heterogeneous behavioral biases, but focus on expectations about aggregate variables in that case. Guerreiro (2023) allows for heterogeneous attention, but focuses on a case where households hold rational expectations about their idiosyncratic shocks. Ilut and Valchev (2023) develop a model of imperfect reasoning and introduce this into an Aiyagari (1994) economy. In contrast to our framework, their households are ex-ante identical and so HtM

[60]The evidence on income forecast errors is more mixed. Souleles (2004) finds evidence of over-optimism in the 1986-1995 Michigan Survey of Consumers (SOC) (see especially his Figure 4 Panel A), using its short panel component to pair one 12-month forecast with a 6-month realization for some respondents. Rozsypal and Schlafmann (ming), using six additional years of SOC data, find that the direction and magnitude of forecast errors vary with income level. d'Haultfoeuille et al. (2021) find that "individuals tend to be right on average about their future earnings", using four-month forecasts and subsequent realizations in the first three waves of the New York Fed's Survey of Consumer Expectations (from 2015). They nevertheless strongly reject rational expectations after accounting for measurement error and aggregate shocks. Caplin et al. (2024) find close alignment between survey income forecasts and administrative data realizations in Denmark.

[61]See, e.g., Woodford (2013), Gabaix (2014), Woodford (2019), Gabaix (2020), Bordalo et al. (2020), Lian (2021), and Boutros (2022).

[62]See, e.g., Farhi and Werning (2019), Auclert et al. (2020), Angeletos and Huo (2021), Laibson et al. (2021), and Pfäuti and Seyrich (2023).

status is driven by adverse idiosyncratic productivity shocks—by bad luck.[63]

A fifth and parallel strand considers (persistent) heterogeneity in reduced-form or presumed-classical preferences. Aguiar et al. (ming) find that allowing for heterogeneity in patience and the elasticity of intertemporal substitution helps match several empirical facts about the behavior of HtM households. They suggest that behavioral factors might provide a potential micro-foundation for their modeling choices. Krueger et al. (2016) and Auclert et al. (2020) introduce permanent heterogeneity in patience and—in the case of Auclert et al. (2020)—in average skills to better match wealth inequality data. Kekre and Lenel (2022) show that heterogeneity in risk aversion can help account for observed heterogeneity in portfolio choice. Kaplan and Violante (2022) show that heterogeneity in risk aversion can produce similar results to heterogeneity in discount factors in terms of HtM shares and MPCs. They also show, however, that allowing for heterogeneity in risk aversion or in discount factors does not solve the standard HANK's "missing middle problem" of producing a wealth distribution that is too polarized. We show that allowing for heterogeneity in overconfidence, in contrast, fills in the missing middle. Furthermore, our micro data does not favor patience or risk aversion alone as an empirically likely key margin of heterogeneity; e.g., their correlations with HtM status are relatively weak compared to cognitive skills and overconfidence, both qualitatively and quantitatively.

Overall, we show that accounting for observed *fundamental* differences between financially constrained and unconstrained consumers is crucial for understanding macroeconomic fluctuations and general equilibrium. This contrasts sharply both with models assuming rational expectations ("RE") and with behavioral models where the only potential deviation from RE regards some aggregate variable. In those classes of models, households become borrowing constrained because they are unlucky, i.e., hit by adverse productivity shocks, and HtM tends to be a relatively transitory state. In our model, households are financially constrained in part because they overestimate their own abilities, leading to a systematic relationship between cognitive skills, overconfidence, and persistent HtM status. Accounting for this relationship turns out to matter greatly for policy as well.

Outline. We detail our data and empirical findings in Section 5.2. Section 5.3 shows how we introduce cognitive skills and overconfidence into HANK models, and Section 5.4 presents our model's stationary equilibrium results. Section 5.5 develops fiscal policy implications and Section 5.6 concludes.

[63] Ilut and Valchev (2023)'s households do not know their optimal policy function and estimate it based on costly (and noisy) deliberation signals. Once households become HtM, they are likely to remain so because they hold excessively high beliefs about their optimal consumption that induce them to dissave and remain at the borrowing constraint. In contrast, HtM households in our setup tend to differ systematically from households away from the borrowing constraint, consistent with what we find in the micro data. Additionally, our model features nominal rigidities and allows for two assets. We also take a step beyond the crucial one of matching key empirical moments by demonstrating use cases for our model: analyzing positive and normative implications for fiscal policy.

5.2 Micro Data and Empirical Results

In this section, we document several new facts regarding consumers' cognitive skills, beliefs about these skills and future financial situations, and how they relate to other forms of persistent heterogeneity and to six measures of hand-to-mouth status. We later use these facts to help discipline and test our model. We show both unweighted and sampling probability-weighted estimates, following Solon et al. (2015).

5.2.1 Data

Our micro data source is the American Life Panel, a long-running online panel that goes to great lengths to obtain a nationally representative sample of U.S. adults.

We measure cognitive skills and overconfidence about cognitive skills using data from the modules in Stango and Zinman (2020) and in Stango and Zinman (2024), henceforth SZ, which elicited behavioral biases and cognitive abilities, together with questions about household financial condition (that we use here to construct some of our measures of HtM status), from the same 845 panelists in two survey rounds administered in 2014 and 2017. The SZ modules sample only working-age adults (aged 18-60 in 2014), which maps well into our model's focus on labor-market productivity. We bring in additional variables—regarding standard measures of HtM status not covered in the SZ modules, and standard measures of consumer sentiment that we use to measure subjective financial condition and expectations thereof—using various other ALP surveys administered from 2010 through 2022. We start by detailing our key variable definitions and prevalences, including comparisons to other work where applicable. We then describe the key micro empirical regularities that shape and discipline our model.

Cognitive skills. We measure cognitive skills for SZ panelists with standard tests for fluid intelligence (McArdle et al. (2007)), numeracy (Banks and Oldfield (2007)), cognitive control/executive function (MacLeod (1991), Miyake and Friedman (2012)), and crystallized intelligence in the form of financial literacy (Lusardi and Mitchell (2014)).[64] We then extract a single common factor (a.k.a. "g" or generalized intelligence) to use as a summary statistic for cognitive skills, as is customary given that various cognitive skills measures are strongly related, both conceptually and empirically (Jensen 1998; Stango and Zinman 2020).[65]

Overconfidence. We measure overconfidence for SZ panelists using the question: "... what you think about your intelligence as it would be measured by a standard test. How do you think your performance would rank, relative to all of the other ALP members who have taken the test?", elicited as an integer percentile. Later in that survey they take a standard

[64] For details on test questions, please see the Data Appendix to Stango and Zinman (2020).

[65] Results are very similar, qualitatively and quantitatively, if we use the first principal component of cognitive skills instead of the first common factor.

15-question "number series" test of fluid intelligence (McArdle et al. (2007)).[66] Respondents are overconfident on average, with 70 percent providing a better-than-average percentile.[67]

We are most interested in *heterogeneity* in overconfidence and measure it in two ways. One is the degree of overconfidence, defined as the self-assessed rank minus the actual rank so that a higher value of this "oc percentile rank" indicates more overconfidence. The second maps into a key model input: the population share of households exhibiting persistent overconfidence. To estimate this input we flag the 38 percent of respondents who are above-median rank in both 2014 and 2017 as "oc in both rounds" (the standard error on this prevalence estimate is 4pp).[68]

We are not aware of any other quantitative estimate of the share of consumers who are persistently overconfident about their ability, or some closely related object, in a plausibly representative national sample of the working-age population. Huffman et al. (2022) estimate that 45 to 48 percent of managers are over-confident about their performance in a repeated high-stakes workplace tournament held by a single employer. Moschini et al. (2023) find widespread over-optimism about college completion among 18 year-olds in the 1997 NLSY. Various theories explain how overconfidence can persist even in the presence of feedback (e.g., Heidhues et al. (2018) or Zimmermann (2020)).

Subjective financial condition forecasts and realizations. We link overconfidence about cognitive skills to consumers' forecasts of their future financial situation. The ALP elicits such forecasts, and subsequent realizations, in many of its survey modules, allowing us to build a panel of 21,586 forecast-realization pairs, provided by 3,467 ALP panelists (including many SZ panelists, as detailed below), across fourteen surveys administered in January and July from July 2010 to January 2016.

The ALP elicits forecasts with a question that has long been used, by the Michigan Survey of Consumers and many other national household surveys across the world, to help measure consumer sentiment (e.g., Souleles (2004)): "... do you think that a year from now you will be better off financially, or worse off, or just about the same as now?". These forecasts are highly correlated with expected income growth in the relatively small number of ALP surveys that also elicit an income forecast (Appendix Table E.1). We measure realizations a year later with "We are interested in how people are getting along financially these days. Would you say that you are better off or worse off financially than you were a year ago?". Both forecasts and forecast errors tilt strongly optimistic in the aggregate,

[66] Number series scores correlate strongly with those from other fluid intelligence tests like IQ and Raven's.

[67] The SZ data provides a second measure of (over)confidence about cognitive skills, regarding absolute performance on the numeracy test, that is strongly correlated with our measure of overconfidence in relative performance (Stango and Zinman (2020), Chapman et al. (2023)). We focus on the relative overconfidence measure because it is more powerful, both statistically (it is more granular in our data) and conceptually (fluid intelligence is linked more strongly to productivity than numeracy is).

[68] Data limitations preclude us from estimating prevalence more precisely, by directly comparing each respondent's forecasted to actual percentile, because the forecast's integer percentile support is much more granular than the 15-question test realization's support.

regardless of the time period in our sample (Appendix Table E.2).[69] Forecast errors are persistent,[70] and there is only modest evidence of learning over relatively long periods of time.[71] Nor is there evidence of substantial overcorrection.[72]

Being especially interested in persistent heterogeneity across consumers, we construct three household-level measures of persistent optimism about financial situations. The first two are indicators equaling one if the proportion of potentially optimistic forecast errors (weakly) exceeds 0.5. The third is the proportion itself. We estimate that 27 to 40 percent of the sample are persistently optimistic in the SZ overlap sample. The SZ sample is key for our subsequent analysis because we have the requisite measures of overconfidence about cognitive skills only for those panelists. We obtain similar estimates of persistent optimism prevalence in the broader ALP sample.

Hand-to-Mouth status. To assess whether someone is (persistently) HtM, we use six different measures of financial constraints. Some of them have been used in previous work, others are new. Two of the six measures are from the two SZ modules. The other four we pull in from other survey modules completed by SZ respondents, so that we can link those additional HtM measures to cognitive skills and overconfidence thereon.

We start by detailing the two HtM measures from the SZ modules. For each of these, we create indicators for whether someone exhibits the symptom of HtM status in both 2014 and 2017. The first measure indicates *severe financial distress*, defined as reporting that any of four events happened in the previous 12 months: forced move, late payments, hunger, or foregone medical care. An estimated 28 or 31 percent of our sample exhibits this indicator in both 2014 and 2017 (for standard errors on these and other estimates of HtM prevalence see Table 5.1 Columns (7) and (8)). Our second measure classifies a household as HtM if its liquid net worth is less than half of total monthly household income. About 40 or 47 percent of our sample exhibits this indicator in both 2014 and 2017. Kaplan and Violante (2022) obtain a similar estimate, of 41 percent, in a snapshot from the 2019 Survey

[69]Appendix Table E.2 shows that forecasts are more than twice as likely to predict improvement (27 to 30 percent of observations) as deterioration (10 to 14 percent of observations). Forecast errors are roughly three times more likely to be in an optimistic than pessimistic direction; to see this, focus on the "same" realization column to allow for the possibility of forecast errors in either direction, and note that an estimated 13 to 18 percent of the sample forecasted better and ended up the same, while only 4 to 7 percent forecasted worse and ended up the same. Our findings are consistent with the evidence of persistent and strong excessive optimism, from several decades of consumer sentiment data across many wealthy countries, in Claus and Nguyen (2023), following Souleles (2004)'s similar findings from 1978-1996 U.S. data. The one counterexample we know of is Hyytinen and Putkuri (2018)'s evidence of aggregate mean-zero forecast errors from Finland.

[70]Appendix Table E.3 shows that 74 percent of consecutive forecast errors are the same (both optimistic, both realistic, or both pessimistic), and that 53 percent of panelists who make an optimistic forecast error in the previous period make the same error in the next period.

[71]Comparing the first to last forecast-realization pair we observe for panelists with multiple pairs, Appendix Table E.4 shows that the accuracy rate increases from 55 to 62 percent and the optimistic slant decreases from 16/21 = 77 percent to 13/18= 72 percent.

[72]Appendix Table E.3 shows that optimists are about 9 times more likely to get better-calibrated than to over-correct with a pessimistic forecast error.

of Consumer Finances.

The third measure of HtM status is indicating strong agreement with the statement "I live from paycheck to paycheck" in a 2012 survey. An estimated 56 or 59 percent of our sample does so. Our fourth measure is closely related and draws on two questions asked in nine COVID-era modules administered May 2020-July 2022. The mean proportion of these modules in which a panelist exhibits paycheck-to-paycheck behavior is about 40 or 44 percent.[73] Our fifth measure indicates whether someone lacks precautionary savings, defined as reporting not having emergency or rainy day funds set aside to cover 3-months of expenses. An estimated 63 or 72 percent of respondents, who completed both surveys where this question was asked, indicate this in either 2012 or 2018.[74] Our sixth measure is based on whether the panelist indicates having difficulty dealing with expense shocks, measured as the proportion of 3 surveys from 2011, 2012, and 2018 where they do not express the highest confidence or certainty that they could cover an unexpected $2,000 need arising in the next month. The mean proportion across panelists is about 51 or 59 percent, as compared to Sergeyev et al. (2024)'s estimate that 54 percent of U.S. households would have difficulty covering an unexpected $2,000 emergency expense in 2022.

Overall, our estimates of HtM prevalence square well with those from prior work. They also suggest that we have measures of financial constraints of varying severity, which will be useful for exploring the robustness of our results below.

5.2.2 Key Correlations

We now use the above variables to estimate the key micro empirical relationships that shape and discipline our model.

[73]For each panelist-survey we define an indicator that =1 if panelists respond "Very difficult" or "Somewhat difficult" to "In the past month, how difficult has it been for you to cover your expenses and pay all your bills?" or, on the followup question "Suppose now you have an emergency expense that costs $400. Based on your current financial situation, how would you pay this expense?" they report one or more expensive options: credit card revolving, small-dollar credit, or that they wouldn't be able to pay for it. For each panelist we then take the ratio of the count of indicators to the count of completed surveys, across the nine modules.

[74]The indicator for lacking precautionary savings is strongly serially correlated within-person across the two surveys, with a tetrachoric correlation of 0.82 (s.e.=0.05) in the sample with nonmissing overconfidence. Unsurprisingly then, correlations are statistically indistinguishable if we define the measure as lacking precautionary saving in both 2012 and 2018. We report results for the either 2012 or 2018 version in the main table, in the interest of showing a measure that indicates relatively high HtM prevalence.

Table 5.1: Pairwise correlations between persistent HtM measures and cognitive skills or persistent overconfidence about skills

	CS rank: cf		1=Oc both rounds		Oc pctile rank		Row variable, unw.	Row variable, w.
	Unw.	Weighted	Unw.	Weighted	Unw.	Weighted	Pop. share	Pop. share
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1=(Severe financial distress)	-0.34	-0.29	0.18	0.27	0.19	0.18	0.28	0.31
s.e.	0.04	0.07	0.06	0.12	0.04	0.08	0.02	0.04
N	841	841	813	813	813	813	813	813
1=(Low net worth)	-0.40	-0.37	0.25	0.20	0.23	0.09	0.40	0.47
s.e.	0.04	0.06	0.06	0.10	0.04	0.07	0.02	0.03
N	788	788	760	760	760	760	760	760
1=(paycheck-to-paycheck c. 2012)	-0.29	-0.50	0.15	0.01	0.15	0.17	0.59	0.56
s.e.	0.07	0.08	0.10	0.24	0.07	0.12	0.03	0.08
N	263	263	255	255	255	255	255	255
paycheck-to-paycheck, COVID	-0.38	-0.28	0.22	0.20	0.30	0.29	0.40	0.44
s.e.	0.02	0.02	0.05	0.09	0.05	0.08	0.04	0.03
N	527	527	516	516	516	516	516	516
1=(Lacks precautionary savings)	-0.30	-0.30	0.11	0.09	0.18	0.19	0.63	0.72
s.e.	0.07	0.12	0.10	0.16	0.07	0.11	0.03	0.04
N	272	272	262	262	262	262	262	262
Difficult covering 2k$ expense	-0.40	-0.43	0.23	0.31	0.22	0.25	0.51	0.59
s.e.	0.04	0.06	0.07	0.09	0.05	0.07	0.04	0.03
N	499	499	485	485	485	485	485	485

Note: CS = cognitive skills, measured as the common factor of four standard tests; OC= overconfidence re: relative performance in a cognitive skills test (see Section 2.1 for details). Weighted estimates use the sampling probability for the last SZ module. In Columns 5 and 6, we use Obviously Related Instrumental Variables to account for measurement error by having the two measurements of o/c rank (taken in 2014 and 2017) instrument for each other (Gillen et al. (2019); Stango and Zinman (2020)). We do not take the same approach to the o/c indicator in Columns 3 and 4, because measurement error-IV does not work well on misclassification error. Fully non-IV correlations estimated using tetrachoric or Pearson. See Section 2.1 for details on HtM measure definitions. The two non-indicator HtM variables are each defined as the proportion of indicators across multiple surveys, so for population share estimates we take the mean of the estimated population shares for each component indicator used in creating that variable.

Empirical strategy

In estimating empirical relationships between variables, we focus on pairwise correlations, for two reasons. One is empirical: pairwise correlations are easier to interpret when all of the variables of interest are correlated with each other; conversely, multivariate estimates are likely subject to confounds from over-controlling and multicollinearity. The other is conceptual: for modeling purposes, we are interested in identifying a proxy for persistent and relatively fundamental consumer heterogeneity (like overconfidence about cognitive skills) that can reproduce key empirical patterns in the aggregate (like patterns of forecast errors and financial constraints). The proxy can be useful, for modeling purposes, whether or not it has a causal relationship with the other variables of interest. We address measurement error in cognitive skills, overconfidence, and other potential sources of fundamental and persistent heterogeneity in decision making by using SZ's repeated measurements as instruments for each other where advisable, following Gillen et al. (2019) and Stango and Zinman (2020).[75]

Cognitive skills and HtM status

As noted at the outset, cognitive skills heterogeneity has been linked to some variables of macroeconomic interest in prior work but not explicitly to HtM status and its persistence within-household over time.[76] Columns (1) and (2) in Table 5.1 take steps towards filling that gap. We estimate unweighted and sampling-probability-weighted correlations between our cognitive skills summary measure and each of our six HtM measures, finding a negative sign on all 12 point estimates. All of them are larger than |0.27| and most have t-stats of |4| or more.

Overconfidence, forecasting, and HtM status

Given that cognitive skills heterogeneity alone is unlikely to help fit the macro data (as we show formally in Section 5.4), we now consider whether overconfidence about cognitive skills is a potential underpinning or proxy for the strong relationship between cognitive skills and persistent HtM documented in Table 5.1 Columns 1 and 2. Indeed, overconfidence in relative performance is the behavioral bias most strongly correlated with cognitive skills out of the 17 biases measured in the SZ data (Stango and Zinman (2020)). Overconfidence could be a key link between cognitive skills and consumer behavior that has been overlooked so far.

Table 5.2 links overconfidence about cognitive skills to over-optimism about one's own future financial situation. We see that persistent optimism about one's own future financial condition—as measured by our two indicators—is about 1.06 to 1.25 times more prevalent

[75]Measurement error IV is advisable for smooth measures but not for discrete ones—the latter are subject to misclassification error that is non-classical.

[76]Recall that 5 of our 6 HtM measures explicitly capture persistence. Because HtM status is so persistent, results on HtM snapshots are similar and we do not report them below.

Table 5.2: Optimistic forecast errors are more prevalent among the overconfident

(Optimist share | overconfident) / (Optimist share | not oc)	Optimism measure	
	1 = (Prop. Opt. FEs > 0.5)	1 = (Prop. Opt. FEs ≥ 0.5)
Unweighted	1.25	1.20
Weighted	1.08	1.06

Note: Sample is the 462 Stango-Zinman panelists who also provide the requisite data, in other ALP modules, to measure at least two potentially optimistic forecast errors. Overconfidence re: relative performance in a cognitive skills test (see Section 2.1 for details). Weighted estimates use the sample probability weights from the last Stango-Zinman module.

among persistently overconfident households than in the rest of the population. In our model calibration, we will use this ratio of relative over-optimism to discipline overconfidence.

Table 5.2 suggests that the strong negative relationship between cognitive skills and HtM status in columns 1 and 2 in Table 5.1 may be due at least in part to overconfidence. Columns 3-6 in Table 5.1 provide empirical support for that conjecture. Here we estimate 24 correlations: (6 HtM measures × 2 overconfidence measures × weighted or unweighted). All 24 point estimates are positively signed, and 17 have t-stats strictly greater than two.[77] Relatedly, Grohmann et al. (2023) find that overconfident participants save less in a lab experiment.

5.2.3 Other sources of fundamental heterogeneity?

Other papers have put forth more classical sources of relatively fundamental heterogeneity as candidates for macro modeling; see e.g., Krueger et al. (2016), Auclert et al. (2020), Aguiar et al. (ming), Kaplan and Violante (2022), and Andreou et al. (2023) on patience, and Kaplan and Violante (2022) and Kekre and Lenel (2022) on risk aversion. But we find that the micro data favors focusing on cognitive skills and overconfidence over patience or risk aversion. Stango and Zinman (2020)'s findings point to cognitive skills heterogeneity as the most likely source or summary statistic for heterogeneity in various behavioral biases, and moreover show that overconfidence in relative performance is the bias that has the strongest correlation with cognitive skills. Here we look directly at relationships between our other key micro variables for macro modeling on the one hand, and patience or risk aversion on the other. We do not find evidence of a robust relationship between those classical decision inputs and persistent over-optimism about financial condition, subject to the caveat that any nulls are imprecisely estimated (Appendix Table E.6). Turning to HtM status, although we do find some evidence of potentially meaningful correlations with patience or risk aversion, overall the relationships are less robustly strong across our six HtM measures than they are with cognitive skills or overconfidence, both statistically

[77] Consistent with Tables 5.1 and 5.2, Table E.5 shows strong correlations between over-optimism about financial condition and HtM status. All 30 point estimates are positive, most have t-stats > 3, and 28 have t-stats larger than 2.

and quantitatively, and patience has a surprising positive correlation with our pre-COVID measure of living paycheck-to-paycheck (Appendix Table E.7). Nor is patience a good proxy for overconfidence (Appendix Table E.8 Columns 1 and 2). Risk aversion might be, but the two different measures of presumed-classical risk aversion in the SZ data have opposite-signed correlations with overconfidence (Appendix Table E.8 Columns 3-6), despite being positively correlated >0.2 with each other.

5.2.4 Summary of results from micro data

To summarize, we find that persistent HtM status decreases strongly with cognitive skills and increases with overconfidence thereon, and that overconfident consumers tend to be persistently too optimistic about their future financial situation. Together with prevalent overconfidence, and the strong negative correlation between cognitive skills and overconfidence found in prior work, these findings suggest that accounting for consumer heterogeneity in cognitive skills and/or overconfidence could be important for understanding macroeconomic fluctuations. We next develop a model to explore this possibility formally and quantitatively.

5.3 Model

We now develop an augmented HANK model, using our new results in Section 5.2, together with consensus estimates of key macro variable moments, to shape and discipline the model. Aside from adding heterogeneity in cognitive skills and overconfidence about these skills, the model is otherwise standard: it features incomplete markets in the spirit of Bewley (1986), Huggett (1993), and Aiyagari (1994), and nominal rigidities in the form of sticky wages. Time is discrete and denoted by $t = 1, 2, \dots$. We first focus on the case in which households can only save in one asset—a liquid bond issued by the government. Later on, we introduce a second asset in the form of illiquid productive capital.

Households. There is a unit mass of households subject to idiosyncratic risk, incomplete markets, and borrowing constraints. We allow for permanent heterogeneity in households' cognitive skills (modelled as productivity) and overconfidence about these cognitive skills (specifically about idiosyncratic productivity).[78] An individual household's productivity of permanent type g in period t are denoted by $\bar{e}_g e_t$, where $\bar{e}_g$ captures permanent differences across groups in average productivity levels, and e_t captures idiosyncratic productivity. The stochastic component e_t follows a Markov process with time-invariant transition matrix $\mathcal{P}$. The process for e_t is the same for all households and the mass of households in state e is always equal to the probability of being in state e in the stationary equilibrium, $p(e)$.

[78] We assume that heterogeneity in cognitive skills and overconfidence is permanent given the results in Stango and Zinman (2024). Consistent with that, Hoffman and Burks (2020) also find, among truckers, that workers' over-optimistic beliefs about their productivity are very persistent.

The problem of an individual household of type g in idiosyncratic state e_t, with beginning-of-period asset holdings b_{t-1}, is given by:

$$V_{g,t}\left(b_{t-1}, e_t\right) = \max_{c_t, b_t} \left\{ \frac{c_t^{1-\gamma}}{1-\gamma} - \frac{n_t^{1+\varphi}}{1+\varphi} + \beta \tilde{\mathbb{E}}_{g,t} V_{g,t+1}\left(b_t, e_{t+1}\right) \right\}$$

subject to

$$c_t + \frac{b_t}{1+r_t} = b_{t-1} + (1-\tau_t) w_t \bar{e}_g e_t n_t \tag{5.1}$$

$$b_t \geq -\underline{b}, \tag{5.2}$$

where c_t denotes consumption, n_t hours worked, r_t the net real interest rate, w_t the real wage, τ_t the income tax rate, and V the value function. We assume a standard CRRA utility function where the parameters γ, φ, and β denote relative risk aversion, the inverse Frisch elasticity of labor supply, and the time discount factor, respectively. These parameters as well as the exogenous borrowing limit $\underline{b}$ are the same for all households and time-invariant.

The expectations operator $\tilde{\mathbb{E}}_{g,t}$ is our key innovation, and we discuss it next.

Cognitive skills and overconfidence. We model heterogeneity in cognitive skill levels as different average productivities $\bar{e}_g$, given the strong (negative) correlation between cognitive skills and income in the data (Stango and Zinman 2020).

All households observe their current productivity $\bar{e}_g e_t$ but overconfident households have biased beliefs about the transition probabilities $p(e_{t+1}|e_t)$. Specifically, overconfident households assign too much probability to reaching (or staying in) relatively high-skill states, and too little probability to reaching (or staying in) relatively low-skill states. This makes overconfident households too optimistic about their expected future productivity, relative to rational households with the same productivity and idiosyncratic risk.

Let $p_{ij} \equiv p(e_{t+1} = e_j | e_t = e_i)$ denote the probability that a household with current idiosyncratic productivity $e_i \in \{e_1, e_2, ..., e_J\}$ reaches productivity $e_j \in \{e_1, e_2, ..., e_J\}$ in the following period, and assume that the productivities are ordered such that $e_1 < e_2 < ... < e_J$. To capture overconfidence with only one additional parameter independent of the number of states, we assume that an overconfident household's *perceived* transition probabilities $\tilde{p}_{ij}$ are given by

$$\tilde{p}_{ij} \equiv \begin{cases} \alpha p_{ij}, & \text{if } i < j \\ \frac{1}{\alpha} p_{ij}, & \text{if } i > j \\ 1 - \sum_{j \neq i} \tilde{p}_{ij}, & \text{if } i = j, \end{cases} \tag{5.3}$$

where the parameter $\alpha \geq 1$ captures overconfidence. If $\alpha > 1$, the household assigns too much weight to reaching a better state (this is the case $i < j$) and too little weight to reaching a worse state ($i > j$). The perceived probability of staying in the same state ($i = j$)

ensures that the probabilities sum to 1.[79] We discuss an alternative modelling approach in Section 5.4.3, where the degree of overconfidence depends on the distance between the states. Note that the rational expectations case is captured by setting $\alpha = 1$ and thus nested in our setup.[80]

An immediate implication is that overconfident households will more often be overly optimistic about their financial situation (specifically income, in the model) compared to rational households, consistent with the empirical findings reported in Section 5.2.2. We will use our empirical estimate of the relative share of optimists among overconfident and rational households from Table 5.2 to calibrate α below (in Section 5.3.1).

Unions. We follow the recent HANK literature and assume that hours worked n_t are determined by union labor demand and that wages are sticky whereas prices are flexible (see especially Auclert et al. (ming), which is based on Erceg et al. (2000)).[81] Each worker provides $n_{k,t}$ hours of work to a continuum of unions indexed by $k \in [0,1]$. Each union aggregates efficient units of work into a union-specific task

$$N_{k,t} = \int \bar{e}_i e_{i,t} n_{i,k,t} di,$$

where i here denotes an individual household carrying its permanent type and in its current idiosyncratic state.

A competitive labor packer then packages these tasks into aggregate employment services according to the CES technology

$$N_t = \left(\int_k N_{k,t}^{\frac{\epsilon-1}{\epsilon}} \right)^{\frac{\epsilon}{\epsilon-1}} \tag{5.4}$$

and sells these services to firms at price w_t.

We model wage stickiness by imposing a quadratic utility cost $\frac{\psi}{2} \int_k \left(\frac{W_{k,t}}{W_{k,t-1}} - 1 \right)^2 dk$ that shows up in the household's utility function. A union sets a common nominal wage $W_{k,t}$ per efficient unit for each of its members.

In doing so, the union trades-off the marginal disutility of working given average hours against the marginal utility of consumption given average consumption. The union then calls upon its members to supply hours. We assume the union ensures that each household supplies the same amount of hours.

[79] We further restrict α such that all perceived transition probabilities lie between 0 and 1. Given a standard calibration for the income process, this restriction is never binding.

[80] Modelling overconfidence as in (5.3) is similar to the way Caballero and Simsek (2019) model optimism about an aggregate state with two possible realizations. In contrast to them, we focus on idiosyncratic states and allow for an arbitrary number of realizations. McClung and Nighswander (2021) introduce belief heterogeneity about idiosyncratic employment transition probabilities into a life-cycle model, but consider only two possible states.

[81] Auclert et al. (2021) and Broer et al. (2020) argue in favor of using sticky wages rather than sticky prices in HANK models.

Firms. A representative firm operates an aggregate production function which is linear in labor input N_t

$$Y_t = N_t, \tag{5.5}$$

to produce total output Y_t. Prices are fully flexible such that the real wage per efficient hour is constant

$$w_t = 1. \tag{5.6}$$

Profits are zero. Since the nominal wage is given by $W_t \equiv w_t P_t = P_t$, we have

$$1 + \pi_t = 1 + \pi_t^w, \tag{5.7}$$

where $\pi_t \equiv \frac{P_t}{P_{t-1}} - 1$ denotes goods price inflation, and $\pi_t^w \equiv \frac{W_t}{W_{t-1}} - 1$ wage inflation.

Fiscal policy. We abstract from government spending and assume that the fiscal authority sets total taxes minus transfers, T_t, following a simple debt feedback rule

$$T_t - \bar{T} = \vartheta \frac{B_t - \bar{B}}{\bar{Y}}, \tag{5.8}$$

where $\bar{T}$, $\bar{B}$ and $\bar{Y}$ denote the stationary equilibrium values of taxes, government debt and output, respectively. Furthermore, the government budget constraint is given by

$$B_t + T_t = (1 + r_t) B_{t-1}. \tag{5.9}$$

Monetary policy. The monetary authority directly controls the real rate r_t and we assume that they keep it constant at its steady state value r. This assumption only matters when we consider aggregate shocks, as we do when examining how overconfident consumers change the effectiveness of temporarily increasing fiscal transfers in Section 5.5.1, .

Equilibrium. Absent aggregate shocks, and given an initial price level P_{-1}, initial nominal wage W_{-1}, initial government debt B_{-1}, and an initial distribution of agents $\Psi_{g,0}(b_{-1}, e_0)$ in each fixed group g, a general equilibrium is a path for prices $\{P_t, W_t, \pi_t, \pi_t^w, r_t, i_t\}$, aggregates $\{Y_t, C_t, N_t, B_t, T_t\}$, individual allocation rules $\{c_{g,t}(b_{t-1}, e_t), b_{g,t}(b_{t-1}, e_t)\}$ and joint distributions of agents $\Psi_{g,t}(b_{t-1}, e_t)$ such that households optimize (given their beliefs), all firms optimize, unions optimize, monetary and fiscal policies follow their rules, and the goods and bond markets clear:

$$\sum_{g,e} \mu_g p(e) \int c_t \Psi_{g,t}(b_{t-1}, e_t) = Y_t \tag{5.10}$$

$$\sum_{g,e} \mu_g p(e) \int b_t \Psi_{g,t}(b_{t-1}, e_t) = B_t, \tag{5.11}$$

where μ_g denotes the mass of agents of type g.

5.3.1 Calibration

One period in the model corresponds to a quarter. We calibrate the standard parameters to values often used in the literature. For the idiosyncratic productivity process, we follow McKay et al. (2016) and set the autocorrelation of e_t to $\rho_e = 0.966$ and the variance to $\sigma_e^2 = 0.033$. We then discretize this process into an eleven-states Markov chain using the Rouwenhorst (1995) method. We set the discount factor, β, to match a steady state real interest rate of 4% (annualized). Risk aversion is set to $\gamma = 2$, the inverse Frisch elasticity to $\varphi = 2$, and the borrowing limit to $\underline{b} = 0$ (as, e.g., in McKay et al. (2016)). We set the average wealth to average annual income ratio to its empirical counterpart of 4.1 (Kaplan and Violante (2022)).

Table 5.3: Persistent overconfidence: prevalence and relationship to income

	Overconfident in both survey rounds?			
	Yes Unweighted	No Unweighted	Yes Weighted	No Weighted
Population share	0.34		0.38	
	(0.02)		(0.04)	
Mean Income	\$51,182	\$79,765	\$42,035	\$77,145
N	817	817	817	817

Note: Standard errors in parentheses. Weighted estimates use the sampling probability for the last SZ module. Income is the sample mean of each panelist's mean income across the two SZ modules.

We set the share of overconfident households to 0.38 (Table 5.3), using the higher estimate from our data in light of Huffman et al. (2022)'s finding of even higher prevalence in a high-stakes workplace tournament.[82] Based on prior work showing strong negative correlations between cognitive skills and overconfidence about those skills (see Ehrlinger et al. (2008); Stango and Zinman (2020) with additional results here in Table E.9), and in the interest of parsimony, we collapse permanent heterogeneity in skills and confidence to two types: overconfident with low skills, and rational with high skills. We normalize the average productivity of the high-skilled and rational households to $\bar{e}_2 = 1$ and set the average skill level of the low-skilled and overconfident households to $\bar{e}_1 = 0.55$, based on our weighted estimates of average income for overconfident vs. rational households in Table 5.3: $0.55 = \frac{42{,}000}{77{,}000}$.

Following equation (5.3), we capture the degree of overconfidence in the overconfident and low-skilled group with one parameter, α. To calibrate α, we target our estimates from Table 5.2 that overconfident households are more likely to have optimistic one-year forecast errors about their financial situation,[83] using a medium value of 1.18 as our target. This

[82]Using our lower estimate of 34% changes our quantitative results only slightly. For example, it changes the share of HtM from 29% to 27% and the average MPC from 0.16 to 0.15.

[83]Note that in the stationary equilibrium of our model a household that is overly optimistic about its future idiosyncratic productivity is also overly optimistic about its future financial situation (defined as labor income plus asset income). The reason is that wages, hours worked, and asset returns are constant and

Table 5.4: Stationary equilibrium calibration

Parameter	Description	Value
R	Steady state real rate (annualized)	4%
γ	Risk aversion	2
φ	Inverse of Frisch elasticity	2
$\underline{b}$	Borrowing limit	0
$\frac{\bar{B}}{4\bar{Y}}$	Average wealth to average income	4.1
Idiosyncratic risk		
ρ_e	Persistence of idiosyncratic risk	0.966
σ_e^2	Variance of idiosyncratic risk	0.033
Permanent heterogeneity		
μ_g	Mass of households	$\{0.38, 0.62\}$
$\bar{e}_g$	Cognitive skills	$\{0.55, 1\}$
α	Degree of overconfidence	2

Note: Calibration summary for our one-asset model using two groups to capture permanent heterogeneity: households in group one have relatively low average skill levels $\bar{e}_1 < \bar{e}_2$ and are overconfident ($\alpha > 1$), group two is relatively high-skilled and has rational expectations ($\alpha = 1$).

results in $\alpha = 2$. Below we consider several alternative parameterizations of heterogeneity in cognitive skills and overconfidence and find similar results. Table 5.4 summarizes our baseline calibration.

5.4 Stationary Equilibrium Predictions

We now consider our model's ability to fit various key moments from macro and micro data, as compared to HANK models that abstract from cognitive skills or belief heterogeneity or both.

5.4.1 Hand-to-Mouth Shares and Average MPCs

We start by considering the effects of permanent heterogeneity in cognitive skills and overconfidence on the share of Hand-to-Mouth (HtM) households and the implied average marginal propensity to consume (MPC) of households.[84]

Table 5.5 compares predictions across four different models: our baseline model with heterogeneity in cognitive skills and overconfidence (*"HANK: CS + OC"*, in Column 1), a standard HANK model (Column 2) with no heterogeneity in permanent productivity levels ($\bar{e}_g = \{1, 1\}$) and full rationality ($\alpha = 1$), a HANK model with permanent heterogeneity in skill levels but full rationality (*"HANK: CS"*, Column 3), and a HANK model with a group of permanently overconfident households but no skill heterogeneity (*"HANK: OC"*, Column

therefore the only possible variation in a household's financial situation comes from changes in idiosyncratic productivity.

[84] Here we define HtM as holding less liquid wealth than half of average monthly income.

4).[85] We start by comparing our model to the standard HANK, and then use the other two models to help unpack the differences.

Table 5.5: MPCs and shares of HtM households across the models.

	HANK: CS + OC (1)	Standard HANK (2)	HANK: CS (3)	HANK: OC (4)
HtM Share	0.29	0.03	0.04	0.28
Avg. MPC	0.16	0.04	0.04	0.18
HtM rational HHs	0.02	0.03	0.03	0.02
Avg. MPC rat. HHs	0.03	0.04	0.04	0.03
HtM OC HHs	-	-	-	0.70
Avg. MPC OC HHs	-	-	-	0.43
HtM rat. HHs Low-Skilled	-	-	0.04	-
Avg. MPC rat. HHs LS	-	-	0.03	-
HtM OC HHs LS	0.74	-	-	-
Avg. MPC OC HHs LS	0.38	-	-	-

Note: MPCs refer to MPCs out of a $500 dollar stimulus check. "HANK: CS + OC" is our baseline model (one-asset, with heterogeneity in cognitive skills and overconfidence); "Standard HANK" denotes a standard one-asset model that abstracts from heterogeneity in skills and overconfidence; "HANK: CS" adds heterogeneity in skills only to Standard HANK; "HANK: OC" adds heterogeneity in overconfidence only to Standard HANK.

Column 2 reproduces the well-documented finding that a standard one-asset HANK model calibrated to match average wealth produces an average MPC and aggregate HtM share that are both far below consensus estimates (Auclert et al. (ming), Kaplan and Violante (2022)). The reason is that rational households have a strong incentive to self-insure themselves against their idiosyncratic risk by accumulating liquid wealth. Thus, with a high enough liquidity supply in the economy, almost no households end up at the borrowing constraint.

In contrast, our model with skill and belief heterogeneity (Column 1) produces an average MPC and a HtM share that are both multiple times larger than in the standard one-asset HANK model. Our predictions align well with consensus estimates, albeit more obviously so for the MPC. For example, Jappelli and Pistaferri (2010) and Havranek and Sokolova (2020) report average MPC estimates in the range of 15-25% over a quarterly time horizon, as compared to our 16%. Our predicted share of HtM households, 0.29, is in the range of our estimated empirical share based on our most conservative definition of HtM status: those with severe financial distress (Table 5.1).

Column 3 shows that skill heterogeneity alone does not drive our model's ability to fit the data better. If we introduce skill heterogeneity but keep all households rational (i.e., well-calibrated about their productivity), the average MPC and the HtM share are very similar to those produced by the standard HANK model. The reason is that a rational household still has a strong incentive to self-insure regardless of its average productivity.

[85] When comparing these four different models, we take the standard approach of recalibrating the discount factor such that all models have the same asset supply and the same steady-state real interest rate (see, e.g., Kaplan and Violante (2022)). The rest of the calibration is the same for all models.

Column 4 shows that our model's allowance for belief heterogeneity drives its improved performance. Specifically, keeping average productivity homogeneous but allowing some households to be overconfident about their future idiosyncratic productivity generates average MPCs and HtM shares that are consistent with the data. The mechanism is that overconfident households overestimate their expected income; i.e., they perceive their income risk to be lower than it actually is. Overconfident households thus accumulate less precautionary savings than rational households facing the same actual income risk. Consequently, and in line with our empirical findings in Section 5.2, overconfident households are much more likely to end up being HtM than rational households in our model (74% of overconfident households are HtM, while only 2% of rational households are). This also results in a high average MPC for the group of low-skilled, overconfident households (0.375, vs. 0.027 for the rational households), driving up the aggregate average MPC. These results are consistent with Bernard (2023)'s empirical finding that a lack of cognitive sophistication is positively correlated with MPCs.

5.4.2 "Missing Middle Problem" and the Top 10% Wealth Share

Standard one-asset HANK models can generate a high average MPC by restricting wealth to be many multiples lower than consensus estimates (Wolf (2021), Kaplan and Violante (2022), Seidl and Seyrich (2023)). This restriction also produces an excessively polarized wealth distribution (Kaplan and Violante 2022). One way to see this "Missing Middle" problem is that median wealth to mean annual earnings is about an order of magnitude smaller in standard HANK models than in the data. We offer further confirmation of this finding by recalibrating the standard HANK model used in Table 5.5 Column 2 to match the average MPC produced by our one-asset model with skill and belief heterogeneity. Matching the average MPC requires setting total wealth to income to 0.7 instead of 4.1, and delivers a median wealth-to-average annual income ratio of 0.2 vs. about 1.5 in the data (Kaplan and Violante 2022).

Our one-asset model with heterogeneity in cognitive skills and overconfidence fills in the missing middle: it predicts a median wealth-to-average annual income ratio of 1.4 that is close to its empirical counterpart of 1.5. Rational households that have experienced several periods of relatively low productivity make up most of the middle of our wealth distribution. Overconfident households tend to be HtM and thus account for most of the bottom, as discussed above. Rational households that have not experienced long spells of bad productivity shocks populate the top of the distribution. Although not targeted, our model predicts that the top 10% of households hold 45% of wealth, as compared to the empirical estimate of 49% (Kaplan and Violante 2022). Overall, our model produces a wealth distribution that matches the data well.

Discount factor heterogeneity. As illustrated by Krueger et al. (2016), Aguiar et al. (ming), or Kaplan and Violante (2022), ex-ante heterogeneity in discount factors β can help the rational model account for some of the MPC patterns observed in the data. Aguiar et al. (ming) suggest that behavioral frictions could provide a microfoundation for the low β of some households. Yet our empirical evidence in Section 5.2 points towards overconfidence and not impatience as having a strong connection to HtM status. Similarly, D'Acunto et al. (2023) find links between other key macro variables and cognitive skills that cannot be explained via heterogeneity in patience.

Besides the empirical evidence, there are also important distinctions from a modeling perspective between heterogeneity in overconfidence and heterogeneity in discount factors. Note first that they are not equivalent, as the following Lemma states.[86]

Lemma 3 *Unless marginal utility is constant across individual states, the model with heterogeneity in overconfidence and the model with heterogeneity in patience are not equivalent.*

The intuition is that overconfidence affects expected marginal utility, which depends on the individual state of a household. In contrast, impatient households have the same lower discount factor independent of their current state. Thus, at the household level, these two models cannot be the same.

At the macro level, it is nevertheless technically possible to produce the same average MPC predicted by our baseline model in a model with discount factor heterogeneity. But this comes at the cost of also producing two unattractive features. First, it requires using the discount factor of the impatient households as a free parameter to match the average MPC. Second, it tends to produce wealth distributions with a missing middle (as shown by Kaplan and Violante (2022)).

We further show that the two models can also differ vastly in their normative implications (Section 5.5.2), highlighting that it matters for the optimal debt level *why* households differ in their savings behavior and HtM status.

5.4.3 Extensions

We now show that our results are robust to: (i) Accounting for the empirical finding that 11% of households are persistently *underconfident*, and (ii) An alternative specification of overconfidence that is state dependent. We then extend the model to incorporate a second, productive asset.

Underconfident households.

Our survey data suggests that 11% of households are persistently underconfident, defined as underestimating their cognitive skills in both survey rounds. We extend our model to

[86] For the proof, see Appendix E.2

account for this by setting $\alpha_{uc} < 1$ for 11% of households and adding a symmetric target to its calibration: we now not only target overconfident households being 1.18 times as likely to be optimistic about their future situations than their rational counterparts, but also underconfident households being 1.18 times as likely to be too pessimistic. We again allow the discount factor to adjust keep the real interest rate at 4% annually.[87]

Incorporating underconfident households actually increases the overall HtM share slightly from 29.2% to 30.0% and the average MPC from 16.3% to 16.7%. In partial equilibrium, one would see effects in the opposite direction, because underconfident households overestimate their precautionary savings motive compared to rational households. This pushes underconfident households to save more than the rational households they are replacing in the model. The underconfident are then slightly less likely to end up HtM (recall that HtM probability is quite low for rational households in any model). In general equilibrium, the added savings demand from underconfident households pushes up the asset price, crowding out savings from the larger mass of households close to the borrowing constraint.

Overall, extending the model by accounting for underconfident households further illustrates how adding heterogeneity in beliefs about skills can help improve model performance in general equilibrium. But given the small share of underconfident households in the data, adding them to our model has only small quantitative effects.

Alternative way of modelling overconfidence.

In our baseline specification of overconfidence (equation (5.3)), the degree of overconfidence is the same for all overconfident households, independent of their current state or skill level. We now allow for dependence of the following form:

$$\tilde{p}_{ij} \equiv \begin{cases} \alpha^{(e_j - e_i)} p_{ij}, & \text{if } i \neq j \\ 1 - \sum_{j \neq i} \tilde{p}_{ij}, & \text{if } i = j. \end{cases} \tag{5.12}$$

As in our baseline specification, when $\alpha > 1$, the transition probabilities of moving upwards ($e_i < e_j$) are overweighted and the probabilities of moving downward are underweighted. Here we posit that these probability distortions are larger for states that are further away from each other.[88]

We again calibrate α to match the empirical finding that overconfident households are about 1.18 times as likely to be overly-optimistic about their future financial situation than rational agents. This implies $\alpha = 2.65$. The predicted average MPC is 0.175 and thus largely unchanged from our baseline estimate of 0.163. The predicted HtM share is now about 10 percentage points higher, at 39.2%, and thus closer to the empirical shares of more expansive definitions of HtM (see Table 5.1).

[87] This requires a discount factor of 0.981 instead of 0.982.

[88] This specification may arise if households' beliefs are more distorted for less-frequent events, such as large changes in their idiosyncratic productivity, than for more-frequent events.

Overconfidence in a Two-Asset Model

Rational HANK models often introduce a second, illiquid asset to match the average MPC while simultaneously matching total wealth in the economy (Kaplan et al. (2018), Kaplan and Violante (2022), Auclert et al. (ming)). This approach seeks to capture illiquid assets that are good long-run savings vehicles but ill-suited for self-insurance purposes. But in order to match high average MPCs, two-asset HANK models typically require a liquidity premium—a return difference between liquid and illiquid assets—that is arguably substantially higher than in the data (Kaplan and Violante (2022)).

We now show that the two-asset version of our model can fit the MPC and wealth data with a substantially lower liquidity premium than required by a standard two-asset HANK model.

Model. Per standard practice, adding an illiquid asset requires enriching the model in two ways. First, households can now save in two assets: a liquid but low-return bond, and illiquid but high-return productive capital. Second, we add capital to the production function.

The household's budget constraint now reads:

$$c_t + \frac{b_t}{1+r_t} + k_t = b_{t-1} + (1+r_t^k)k_{t-1} + (1-\tau_t)w_t\bar{e}_g e_t n_t, \tag{5.13}$$

where k denotes the illiquid asset of the household and r^k is its net return. Capital depreciates at rate δ and depreciated capital has to be replaced for maintenance. We follow Bayer et al. (ming) and assume that households make their savings and portfolio choices between liquid bonds and illiquid capital in light of a capital market friction: participation in the capital market is random and i.i.d. in the sense that only a fraction λ of households can adjust their capital holdings in a given period. Households not participating in the capital market in a given period ($k_t = k_{t-1}$) still obtain the return on their illiquid asset holdings and can adjust their bond holdings. We further assume that holdings of both assets must be non-negative:

$$b_t, k_t \geq 0.$$

A representative firm operates a Cobb-Douglas production function using capital (K) and labor (N) as input factors:

$$Y_t = K_{t-1}^{\chi} N_t^{1-\chi}, \tag{5.14}$$

where χ denotes the capital share in production.

In addition to the equilibrium conditions in Section 5.3, now the capital market must clear:

$$\sum_{g,e} \mu_g p(e) \int k_t \Psi_{g,t}\left(k_{t-1}, e_t\right) = K_t. \tag{5.15}$$

Calibration. We maintain the same values for each of the parameters that also appear in our baseline model (except for the discount factor). Table 5.6 shows our calibration of the additional parameters and the discount factor. We set the capital share to $\chi = 0.318$ and the quarterly depreciation rate to $\delta = 0.0175$ as in Bayer et al. (ming). We then use the per-period capital market participation probability λ and the discount factor β to jointly target the average wealth-to-annual income ratio of 4.1 and the liquid asset-to-annual income of 0.2 as in Kaplan and Violante (2022).

Table 5.6: Calibration two-asset model

Parameter	Description	Value
χ	Capital share	0.318
δ	Depreciation rate	0.0175
λ	Capital market participation rate	0.37
β	Discount factor	0.992

Note: The table shows the values for the additional parameters in our two-asset model and the discount factor. All other parameters stay the same as in our baseline model.

Stationary Equilibrium Results. Table 5.7 shows the influence of overconfident households on the stationary equilibrium (Column 1). We start by explaining the mechanisms underlying our results, and then compare the empirical fit of our model to standard models.

The share of HtM households is now 0.38, as compared to 0.29 in our baseline model, because the illiquid asset's higher return induces some savers to substitute from the liquid asset. This is mostly driven by "wealthy HtM" households who would not be HtM in a one-asset model and now choose to save only in the illiquid asset (Kaplan et al. (2018)). In contrast, the average MPC increases by only about one percentage point here relative to our baseline model, to 0.171, indicating that in the two-asset model, on average, constrained households are not as far off their Euler equation and thus spend less out of a \$500 windfall. There is again a stark difference between the behavior of rational and overconfident households in our model. Rational households accumulate liquid assets to self-insure before saving in the illiquid asset. Overconfident households remain much more likely to be HtM (77% vs. 14%) because they foresee little value in accumulating a liquid buffer stock and hence prioritize the illiquid asset's higher return if they do save.

Table 5.7: MPCs and liquidity spread across two-asset models.

	Two-asset HANK w overconfidence	Rational two-asset HANK	
	(1)	(2)	(3)
		Calibrated as (1)	Recalibrated
HtM	0.38	0.23	0.27
Avg. MPC	0.17	0.06	0.15
return gap (annualized)	2.3%	4.4%	9.3%
HtM rat. HHs	0.14	0.23	0.27
Avg. MPC rat. HHs	0.04	0.06	0.15
HtM OC HHs	0.77	-	-
Avg. MPC OC HHs	0.39	-	-

Note: MPCs refer to MPCs out of a stimulus check of $500. The model in Column 3 is recalibrated to produce an average MPC of 0.15.

Column 2 presents a standard two-asset HANK model for comparison, keeping all the parameters the same as in our model except for recalibrating β to target the mean wealth-to-annual income ratio of 4.1.[89] Unlike our model, this produces a HtM share of 0.23 and average quarterly MPC of 0.06 that are substantially below the lower end of the consensus ranges of empirical estimates.[90] Targeting an average MPC at the lower end, e.g. 0.15, requires an annualized return gap of 9.3% (Column 3).[91] Our model produces a much lower return gap of 2.3% because overconfident households underestimate precautionary savings needs and thus require a much smaller premium on illiquid assets, thereby driving demand for the illiquid asset up and its return down.

Given empirical estimates of the return gap in the ballpark of 5% (see, e.g., Jordà et al. (2019)), it may seem at first glance that our two-asset model undershoots substantially. But both our model and standard HANK abstract from aggregate risk. Accounting for aggregate risk would likely push our estimated risk premium closer to the data and a standard HANK model's estimate even farther away from it, in the case where standard HANK targets an empirically realistic average MPC as in Column 3.

[89] This requires quarterly $\beta = 0.989$, as compared to 0.992 in our model.

[90] Compared to our model, the average MPC in the rational model is also lower conditional on the share of HtM households, because overconfident households have higher MPCs conditional on their current state given their lower perceived income risk.

[91] In targeting the quarterly average MPC of 0.15 we set $\beta = 0.9805$, $\lambda = 0.15$, and $\delta = 0.00875$.

5.5 Policy Implications of the Systematic Relationship between Cognitive Skills, Overconfidence, and HtM Status

We now show that heterogeneous overconfidence matters for the design and effectiveness of fiscal policy tools seeking to stimulate and/or insure consumption. There are two key mechanisms. First, our model produces more higher-income HtM households, as overconfidence is a key predictor of HtM status even conditional on income. This has implications for the effectiveness of income-targeted transfer payments on stimulating private consumption. Second, overconfident households undervalue self-insurance and hence are less responsive to changes in precautionary savings incentives. This dampens crowdout when the government provides insurance (we consider a minimum income benefit as an example), but makes it more difficult to induce households in the neighborhood of the borrowing constraint to self-insure when the government provides liquidity (through higher public debt levels).

5.5.1 The distribution of HtM households and targeted transfers

We start by considering an income-targeted transfer that is not anticipated by consumers. Given the difficulty of directly identifying empirical evidence on the general-equilibrium effects of transfer policies, they are generally evaluated using models—models that match the observed average MPC (see e.g., Kaplan and Violante (2014), Wolf (2021)). But for targeted transfers it is the MPC of transfer recipients that matters most.

As such, matching the HtM-income distribution is important for accurately evaluating the stimulative efficacy of targeted transfer policies (Figure 5.1). The figure's gray hashes show the relationship between HtM status and income found in our microdata (with plus/minus two standard errors also in grey).[92] The standard one-asset HANK model with high wealth, depicted by the black diamonds, unsurprisingly underestimates the HtM shares at all income levels: almost everyone saves their way out of low wealth for precautionary reasons, even at low incomes. The standard HANK model recalibrated to generate the same average MPC as our baseline model, depicted by the red squares, produces more HtM households but far too many of them are low-income. Low income predicts HtM status counterfactually strongly in standard models because households become HtM solely due to bad luck. Our model (depicted by the blue dots) also has overconfident households choosing to be HtM throughout the bottom three income quartiles, thereby better matching empirical estimates.

[92] We use our "severe financial stress" empirical HtM measure for this comparison because it yields the same aggregate HtM share as our baseline model, thereby giving the model an opportunity to match the HtM shares along the income distribution. Although the levels of our different HtM measures differ quite substantially, their relative steepness along the income distribution are similar. Appendix Figure E.11 shows this for our HtM measure based on liquid net worth-to-income.

Figure 5.1: HtM shares along the income distribution

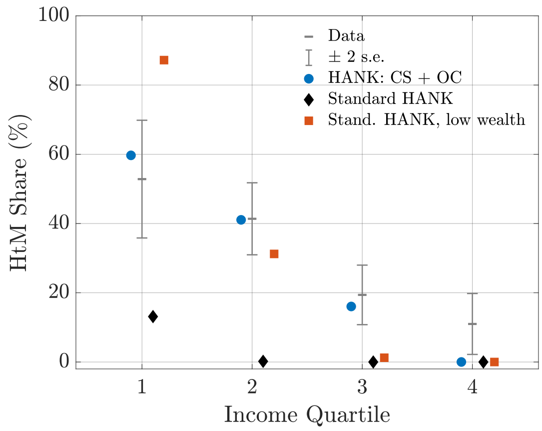

Note: "HANK: CS + OC" is our baseline model. "Standard HANK, low wealth" is the standard HANK model recalibrated to match the average MPC of our baseline model. "Data" shows our prevalence estimates for the severe financial distress HtM measure.

Our model's more realistic depiction of the HtM-income gradient implies that transfers targeted to low-income households are less effective at stimulating consumption than the standard HANK model with the same average MPC would imply, because income is a much weaker predictor of MPCs in our model. Consider a surprise lump-sum transfer to each household in the bottom income quartile, in an aggregate amount of 1 percent of steady-state output on impact, following an AR(1) process with a persistence parameter of 0.8, and financed in the short-run by higher debt which is then slowly repaid with higher taxes. Figure 5.2(a) shows the exogenous path of the transfers, and 5.2(b) shows output, in terms of percentage deviations from steady-state output. Our model (dashed blue line) predicts an output response that is less than half as strong than in the standard HANK model with the same average MPC (dotted orange line): the standard HANK model implies a transfer multiplier of about 2.4 on impact whereas our model implies 0.9.[93] These comparisons highlight that the average MPC is not a sufficient statistic for analyzing targeted stimulus: the distribution of MPCs is important as well.

A second channel further weakens the effectiveness of targeted transfers in our model: muted relaxation of the precautionary saving motive. The persistence in the transfer payments provides some insurance, temporarily decreasing households' precautionary savings motive. This motive is prevalent and strong in the standard model with classically rational households, further increasing spending and total output.[94] But overconfident households undervalue the insurance because they underestimate the likelihood of being income-eligible

[93]The low-MPC standard HANK model (solid black line) produces a multiplier of about 0.5, due to its low MPC across all income groups.

[94]See e.g. Bayer et al. (2020)'s analysis of targeted transfers in a rational HANK model where the relaxation of households' precautionary savings is an important contributor to high multipliers. Kekre (2021) and Dengler and Gehrke (2023) find similar results for temporary increases in unemployment benefits and "short-term work", both of which can be understood as targeted transfers although they are not lump-sum and thus additionally have distortionary effects. Beraja and Zorzi (2024) analyze potential size-dependency for stimulus transfers.

Figure 5.2: Targeted Transfer Shocks

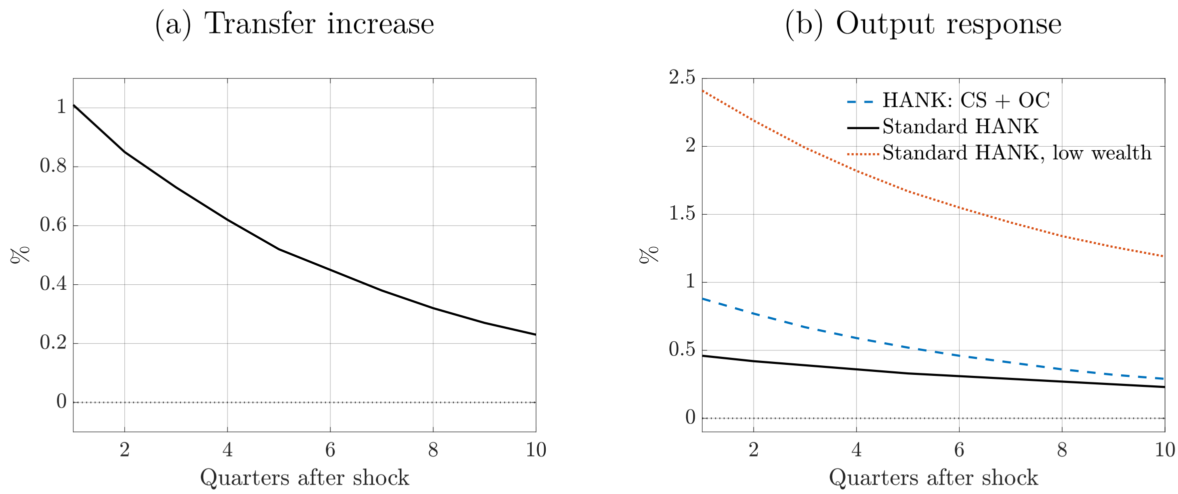

Note: This figure shows the effects of a positive transfer shock (left panel) on total output (right panel). Both are expressed in percentage deviations from steady state output. "HANK: CS + OC" is our baseline model. "Standard HANK, low wealth" is the standard HANK model recalibrated to match the average MPC of our baseline model..

to exercise the insurance option in the future. They thus *perceive* their precautionary savings motive to be less relaxed than rational households would, and as such do not increase their spending as much.[95] This second channel mirrors recent empirical evidence that cognitive constraints can limit the effectiveness of macro policies designed to induce behavior change through incentive changes (D'Acunto et al. 2023).

5.5.2 Precautionary Savings Behavior and Fiscal Insurance Policies

Accounting for the muted responsiveness of overconfident households to changes in precautionary savings incentives is even more crucial when modeling the impact of fiscal policies focused on insurance provision. We now consider two such policies: minimum income benefits as a form of public insurance, and government liquidity provision that reduces the cost of private insurance.

Minimum income benefits as public insurance

We start by analyzing the effects of introducing minimum income benefits (MIB) that provide some public insurance against households' income risk. Following Bayer et al. (2023), we model MIB as a transfer $tr_{i,t}$ to household i contingent on the household's pre-tax labor income $w_t n_{i,t} e_{i,t}$ falling short of some threshold level:

$$tr_{i,t} = \max\{0, a_1 \bar{y} - a_2 w_t n_{i,t} e_{i,t}\},$$

[95] The relaxation of the precautionary savings motive is also an important driver in the standard HANK model with low average MPCs (black-solid line in Figure 5.2). But the MPCs are so low in that model, across all income quartiles, that it still predicts a smaller effect on aggregate output than our model.

where $\bar{y}$ is the median income in the stationary equilibrium and $0 \leq a_1, a_2 \leq 1$. Transfers thus decrease in individual income at the withdrawal rate a_2 and no transfers are paid to households whose labor income satisfies $w_t n_{i,t} e_{i,t} \geq \frac{a_1}{a_2}\bar{y}$. Following Bayer et al. (2023), we set $a_1 = 0.5$ and $a_2 = 0.8$. and assume for simplicity that these transfers do not distort labor supply.

Total government transfer payments are then:

$$Tr_t = \mathbb{E}_t tr_{it},$$

where the expectation operator is the cross-sectional average. These transfers are financed via labor-income taxes.

Table 5.8: Effects of introducing public insurance

	HANK: CS + OC (1)	Standard HANK (2)	Standard HANK, low wealth (3)
HtM Share	0.29	0.03	0.30
Avg. MPC	0.16	0.04	0.16
Bottom50W	2.7%	12.8%	3.0%
Real rate	4%	4%	4%
HtM Share with MIB	0.32	0.09	0.40
Avg. MPC with MIB	0.15	0.06	0.26
Bottom50W with MIB	1.6%	9.2%	1.3%
Real rate with MIB	5.0%	5.5%	6.9%

Note: MPCs refer to MPCs out of a \$500 dollar stimulus check. "HANK: CS + OC" is our baseline model (one-asset, with heterogeneity in cognitive skills and overconfidence), "Standard HANK" denotes a standard one-asset model, in which we abstract from heterogeneity in skills and overconfidence, "Standard HANK low wealth" is the same HANK model but with restricted liquidity to match the average MPC of "HANK: CS + OC". "... with MIB" refers to the stationary equilibrium in the models with public insurance via minimum income benefits (MIB).

Table 5.8 compares the stationary equilibrium effects of MIB on the average MPC and HtM share in our baseline model (Column 1) to a standard rational one-asset HANK model (Column 2). We also again consider a standard HANK model in which we reduce the amount of wealth such that it produces the same average MPC in the absence of transfers as our model does (Column 3).

In the two standard models, targeted transfers crowd-out self-insurance precautionary savings in the stationary equilibrium quite strongly. Households correctly forecast the probability of a bad productivity shock and thus internalize the insurance value of receiving a transfer in that state, reducing their precautionary savings accordingly. This increases the average MPC by more than 50% in either standard model, and the HtM share also increases substantially (by 6pp from the low base in Column 2, and by 10pp on the base of 30 in Column 3). Crowd-out is also reflected by the large increase in the equilibrium real interest rate from 4% to 6.9%. This higher rate is required to induce non-HtM households to hold the liquidity foregone by those moving to the borrowing constraint in response to the policy. Overall then, under standard HANK, introducing minimum income benefits as social insurance produces an economy with substantially higher interest rates, less precautionary savings, more HtM households, and a higher average MPC.

In our model, crowd-out and its concomitant effects are dampened because overconfident households underpredict their probability of reaching a low-productivity state in which they receive a transfer. The average MPC even slightly decreases from 0.163 to 0.151,[96] while the share of HtM households only mildly increases from 29.2% to 32.1%. The real interest rate increase is also substantially smaller, rising only to 5.0%.

Liquidity Provision and the Optimal Public Debt Level

Fiscal policy can also facilitate private insurance, by issuing more government debt (e.g., Woodford (1990)). More debt increases the supply of liquid assets and thus of self-insurance possibilities for households. But this increase in liquidity supply has muted effects in our model compared to the rational HANK model. Figure 5.3a shows the share of HtM households, and 3b the share of wealth held by the poorest 50% of households, as a function of the government debt level in steady state.[97]

The solid black lines in Figure 5.3 show that in the standard, rational HANK model, the provision of liquidity drives down the share of HtM, and increases the wealth share of the bottom 50%, quite effectively. Households at or near the borrowing constraint have the strongest incentive to self-insure by saving in liquid assets and respond strongly as the price of liquidity falls. This drives down their HtM likelihood such that for relatively high public debt levels, almost no households are borrowing constrained.

The dashed blue lines in Figure 3 illustrate the much weaker household response to liquidity provision in our model. The share of HtM households has a relatively flat slope with respect to debt supply, and it plateaus well above zero; e.g., it is about 0.29 at a debt-to-GDP ratio of 4, compared to nearly zero in the standard model. The bottom 50% wealth slope is remarkably flat, reaching only about a 3% share at a debt-to-GDP ratio of 4 compared to about 13% in the standard model. Even when liquidity is abundant, overconfident households do not tend to save themselves out of being liquidity constrained because they still perceive the liquid asset price as too high compared to their underestimated income risk.

The relative unresponsiveness of households at or close to the borrowing constraint in our model also has implications for the optimal amount of government debt. A social planner weighs the benefits of smoother household consumption (from cheaper self-insurance)

[96] There are two opposing effects of the introduction of MIB on the average MPC: on the one hand, the effective lower income risk reduces households' MPC conditional on their individual state. On the other hand, there are more households in individual states with higher MPCs as MIB crowd out precautionary savings. In the rational models, the latter dominates whereas in our baseline model, the first effect dominates because minimum income benefits only mildly crowd out households' precautionary savings.

[97] When varying the supply of government debt, we fix the discount factor β as calibrated in Table 5.4 and let the interest rate adjust to clear the bond market.

Figure 5.3: The Implications of Higher Government Debt

(a) HtM Share (b) Bottom 50% Wealth Share

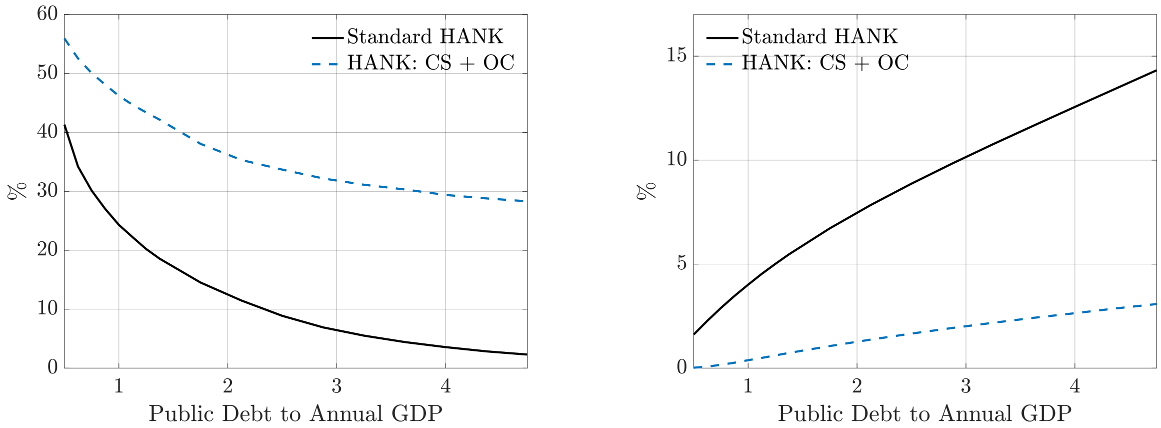

Note: This figure shows the share of HtM households in panel (a) and the wealth share of the bottom 50% of households in panel (b) for varying degrees of average government debt to average earnings ratios (horizontal axis). The black-solid lines show the case for the one-asset standard HANK model that abstracts from permanent heterogeneity in cognitive skills and overconfidence, and the blue-dashed lines show the case for our baseline HANK model featuring permanent heterogeneity in cognitive skills and overconfidence.

vs. the costs of the distortionary taxes required to finance the government's additional interest rate payments. We evaluate this trade-off in both models using a utilitarian social welfare function that seeks to maximize the average expected discounted lifetime utility of households.[98]

Figure 5.4: Public Debt and Social Welfare

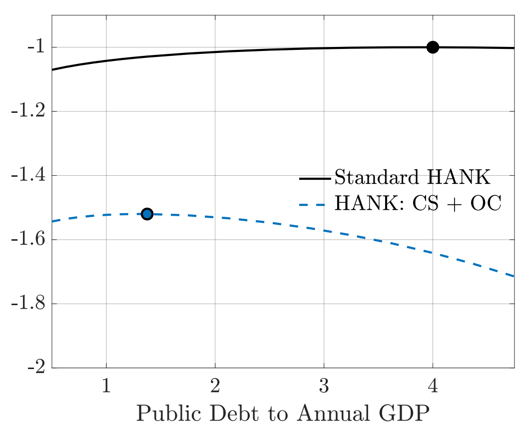

Note: This figure shows average welfare, defined as average expected discounted lifetime utility, as a function of government debt. Dots show the welfare-maximizing amount of government debt for our baseline model (blue-dashed lines) and its rational counterpart (black-solid line). The y-axis shows (normalized) average expected lifetime utility, and the x-axis shows (Public debt outstanding)/(Annual GDP), $\frac{B}{4Y}$. For readability, we normalize welfare such that the highest level of welfare in the model with rational expectations is normalized to -1.

Figure 5.4 shows that average welfare peaks at a much lower debt level in our model compared to the standard one-asset HANK model: optimal debt is about 135% of annual

[98]Such an objective function takes into account aggregate efficiency, risk-sharing, and intertemporal-sharing (Dávila and Schaab 2023). The expectations over the individual lifetime utilities in the social welfare function are assumed to be rational, in the spirit of what Benigno and Paciello (2014) call "paternalistic".

GDP, compared to about 400% in the standard HANK model. Since overconfident households underestimate their income risk and therefore have a dampened response to the liquidity supply increase even when they are at or close to the borrowing constraint, the very households that the social planner would like to save more are the least responsive ones. This diminishes the social benefit of higher government debt compared to the standard model. Even though we abstract from many important channels here—and therefore, our quantitative estimates should be interpreted with caution—the mechanism through which heterogeneity in overconfidence reduces the optimal debt level likely holds in richer models as well.[99]

Analyzing the optimal debt level also highlights the importance of accounting for *why* households differ in their savings behavior and HtM status. For example, our model and a model with heterogeneity in discount factors produce very different optimal debt levels, even when we consider the discount factor heterogeneity model that produces the same average MPC at our baseline wealth-to-income ratio of 4.1. In the model with discount factor heterogeneity, the optimal debt level is 2.5 times as high as in our baseline model because the households who benefit more from government liquidity provision (those with higher discount factors, because they value precautionary savings more) also get de facto higher social welfare weights in a utilitarian welfare function (because their future utility is discounted less). As such, accounting for the strong empirical relationships between overconfidence, savings behavior and HtM status in Table 5.1, rather than relying on heterogeneity in patience (and its weaker empirical links to HtM status in Table E.7), can matter greatly for optimal policy.

5.6 Conclusion

We analyze implications of heterogeneity in cognitive skills and self-perceptions thereof for households' savings behavior and financial situations, macroeconomic fluctuations, and fiscal policy. We start with U.S. micro data and find that lower-skilled households systematically overestimate their skills and are persistently overly optimistic about their future financial situations. They are also substantially more likely to be persistently HtM.

Guided by these findings, we then introduce persistent heterogeneity in skills and overconfidence into a HANK model and uncover a systematic reason why many households are persistently HtM: "bad decisions", not just "bad luck". Accounting for this reason, in the form of overconfidence about future productivity, resolves heretofore seemingly intrinsic tensions in HANK models. Unlike other models, our one-asset HANK model can simultaneously match consensus estimates of both the average MPC and the average wealth

[99] In a robustness exercise, we analyze the optimal debt level in our two-asset model and its rational counterpart. Overconfidence again reduces the optimal debt level significantly, although for both models the level of optimal debt is lower than in the respective one-asset models due to crowdout of productive capital. See e.g., Aiyagari and McGrattan (1998), Davila et al. (2012), Angeletos et al. (2023), or Woodford (1990) for analyses of optimal public liquidity provision.

level. Our model also matches the income-HtM distribution whereas the rational model does not. Our two-asset HANK model matches the data with a lower, and perhaps more empirically realistic, liquidity premium than required in other models. It turns out that our key innovation is the overconfidence of low-skilled households rather than their lower productivity level. Thus, our model requires only one additional parameter—the degree of overconfidence of low-skilled households, as disciplined by our empirical findings—to substantially improve the empirical fit of existing HANK models.

We also show that accounting for the underlying reason why some households are persistently financially constrained matters greatly for fiscal policies. This is particularly pronounced for policies that affect the precautionary savings incentives of households, because overconfident households undervalue insurance and thus have muted responses to changes in such incentives. It also matters for income-targeted transfers, because in our model—as in the micro data—income is much less strongly correlated with HtM status and hence much less of a summary statistic for the MPC than in standard models.

One consideration for future work on normative questions—we mostly consider positive ones in this paper—is whether overconfidence may not be all bad, from a welfare perspective (as in, e.g., Brunnermeier and Parker (2005)). If it is not all bad, quantitative welfare modeling might seek to account for the benefits. Regardless, our finding that overconfidence correlates strongly with persistent and severe financial distress suggests important costs—costs that might be amplified by financial stress (Sergeyev et al. (2024)).

We also stop short of examining different combinations of macroeconomic policies in the presence of permanent heterogeneity across households—but our model provides a framework for doing so going forward. Consideration of monetary policy, and fuller consideration of fiscal policy, likely will require accounting for an additional source of heterogeneity: beliefs about aggregate variables. Some recent papers find empirical links between heterogeneity in expectations about such variables and cognitive skills (D'Acunto et al. (2019b), D'Acunto et al. (2023)). Modeling such links should be a fruitful new line of inquiry.

Appendix A

Appendix for Chapter 1

A.1 A Climate-HANK model

The model in the paper is based on the two-asset, medium-scale HANK model in Bayer et al. (2020b). I extend the model to cover carbon dioxide emissions in production and in household consumption.

The economy consists of a firm sector and a household sector. The firm sector comprises (a) perfectly competitive intermediate goods producers, who produce intermediate goods using capital, labor, and carbon dioxide; (b) final goods producers that face monopolistic competition when selling differentiated final goods, in turn, produced on the basis of homogeneous intermediate inputs; (c) producers of capital goods that turn consumption goods into capital subject to adjustment costs; (d) labor packers that produce labor services combining differentiated labor from (e) unions that differentiate raw labor rented out from households. Price setting for the final goods, as well as wage setting by unions, is subject to a pricing friction à la Calvo (1983).

Households consume a bundle that consists of produced goods and carbon dioxide directly. Households earn income from supplying (raw) labor and capital to the national labor and the national capital markets and from owning firms in their respective country. Households absorb all rents that stem from the market power of unions and final good producers, and decreasing returns to scale in capital goods production.

There is a monetary authority and a fiscal authority. The fiscal authority levies taxes on labor income and profits, issues bonds, pays transfers, sells carbon certificates, and adjusts taxes to stabilize the level of outstanding debt in the long run. Public debt is risk-free and, in turn, determined by monetary policy by means of a simple interest rate feedback rule.

A.1.1 Households

The household sector is subdivided into two types of agents: workers and entrepreneurs. The transition between both types is stochastic. Both rent out physical capital, but only workers supply labor. The efficiency of a worker's labor evolves randomly exposing households to labor-income risk. Entrepreneurs do not work but earn all pure rents in the economy except for the rents of unions which are equally distributed across workers.

All households self-insure against the income risks they face by saving in a liquid nominal asset (bonds) and a less liquid asset (capital). Trading illiquid assets is subject to random participation in the capital market. To be specific, there is a continuum of ex-ante identical households of measure 1, indexed by i. Households are infinitely lived, have time-separable preferences with time discount factor β, and derive felicity from consumption and leisure. Total consumption c_{it} consists of carbon dioxide, E_{it}^C, and the physical consumption good c_{it}^P. Households obtain income from supplying labor, n_{it}, from renting out capital, k_{it}, and from earning interest on bonds, b_{it}, and potentially from profits or union transfers. Households pay taxes on labor and profit income and receive minimum income benefits as well as other transfers.

Productivity, labor supply, and labor income

A household's gross labor income $w_t n_{it} h_{it}$ is composed of the aggregate wage rate on raw labor, w_t, the household's hours worked, n_{it}, and its idiosyncratic labor productivity, h_{it}. I assume that productivity evolves according to a log-AR(1) process with time-varying volatility and a fixed probability of transition between the worker and the entrepreneur state:

$$\tilde{h}_{it} = \begin{cases} \exp(\rho_h \log \tilde{h}_{it-1} + \epsilon_{it}^h) & \text{with probability } 1-\zeta \text{ if } h_{it-1} \neq 0, \\ 1 & \text{with probability } \iota \text{ if } h_{it-1} = 0, \\ 0 & \text{else.} \end{cases} \tag{A.1}$$

with individual productivity $h_{it} = \frac{\tilde{h}_{it}}{\int \tilde{h}_{it} di}$ such that $\tilde{h}_{it}$ is scaled by its cross-sectional average, $\tilde{h}_{it} di$, to make sure that average worker productivity is constant. The shocks ϵ_{it}^h to productivity are normally distributed with variance $\sigma_{h,t}^2$. With probability ζ households become entrepreneurs ($h = 0$). With probability ι an entrepreneur returns to the labor force with median productivity. An entrepreneur obtains a share of the pure rents (aside from union rents), Π_t^F, in the economy (from monopolistic competition in the goods sector and the creation of capital). I assume that the claim to the pure rent cannot be traded as an asset. Union rents, Π_t^U are distributed lump sum across workers, leading to labor-income compression. For tractability, I assume union profits to be taxed at a fixed rate independent of the recipient's labor income.

With respect to leisure and consumption, households have Greenwood et al. (1988) (GHH) preferences and maximize the discounted sum of felicity:

$$E_0 \max_{\{c_{it}, n_{it}\}} \sum_{t=0}^{\infty} \beta^t u[c_{it} - G(h_{it}, n_{it})] \tag{A.2}$$

Total consumption c_{it} of household i at time t consists of carbon dioxide E^C_{it} and the physical consumption good c^P_{it}, again combined in a CES aggregator:

$$c_{it} = \left(\left(1 - a^C_{it}\right)^{\frac{1}{\sigma_C}} c^P_{it}{}^{\frac{\sigma_C - 1}{\sigma_C}} + a^C_{it}{}^{\frac{1}{\sigma_C}} \left(E^C_{it}\right)^{\frac{\sigma_C - 1}{\sigma_C}} \right)^{\frac{\sigma_C}{\sigma_C - 1}}. \tag{A.3}$$

Here σ_C represents the elasticity of substitution in consumption, which determines how much utility the household loses by substituting carbon dioxide for physical consumption goods. a^C_{it} determines the share of the carbon dioxide in the consumption good. The parameter follows a Markov chain to capture households with relatively high carbon dioxide intensity as well as households with relatively low carbon dioxide intensity. The switching probability $\rho(h, a^C)$ from one type to the other is a function of the current productivity level, h, and the current carbon dioxide intensity, a^C. I specify

$$\rho(h, a^C) = \bar{\rho} + (\mathbb{I}_{a^C = a^C_H} - \mathbb{I}_{a^C = a^C_L})A(h) + \mathbb{I}_{a^C = a^C_L} B,$$

where A is a linear function of the human capital quintile. With higher human capital the household is more likely to remain type low and more likely become type low. B is a constant that captures that it is in general more likely to remain type low.

The maximization is subject to the budget constraints described further below. The felicity function u exhibits a constant relative risk aversion (CRRA) with risk aversion parameter $\xi > 0$,

$$u(x_{it}) = \frac{1}{1 - \xi} x_{it}^{1-\xi}, \tag{A.4}$$

where $x_{it} = c_{it} - G(h_{it}, n_{it})$ is household i's composite demand for (carbon dioxide and physical composite) goods consumption c_{it} and leisure and G measures the dis-utility from work.

The household's labor income gets taxed at rate τ_t, such that its net labor income, expressed in physical consumption units (i.e. without carbon dioxide consumption), is given by

$$y_{it} := (1 - \tau_t) w_t h_{it} n_{it}, \tag{A.5}$$

where w_t is the aggregate real wage rate (in physical consumption units). Given net labor income, the first-order condition for labor supply is

$$\frac{\partial G(h_{it}, n_{it})}{\partial n_{it}} = (1 - \tau_t) \frac{w_t}{p^c_t(a^C_{it})} h_{it} = \frac{y_{it}}{n_{it}} / p^c_t(a^C_{it}). \tag{A.6}$$

Here $p^c_t(a^C_{it})$ is the cost in terms of physical goods at which household i buys its carbon dioxide-physical consumption bundle. This price depends on the carbon intensity of the household and is given by

$$p_t^c(a_{it}^C) = \left[(1 - a_{it}^C) + a_{it}^C(p_t^E - \tau_t^E)^{1-\sigma_C}\right]^{\frac{1}{1-\sigma_C}}.$$

Assuming that G has a constant elasticity w.r.t. n, $\frac{\partial G(h_{it},n_{it})}{n_{it}} = (1+\gamma)\frac{G(h_{it},n_{it})}{n_{it}}$ with $\gamma > 0$, I can simplify the expression for the composite consumption good, x_{it}, making use of this first-order condition (C.21), and substitute $G(h_{it}, n_{it})$ out of the individual planning problem:

$$x_{it} = c_{it} - G(h_{it}, n_{it}) = c_{it} - \frac{1}{1+\gamma} y_{it}/p_t^c(a_{it}^C). \tag{A.7}$$

When the Frisch elasticity of labor supply is constant and the tax schedule has the form (C.20), the dis-utility of labor is always a fraction of labor income and constant across households. Therefore, in both the household's budget constraint and felicity function, only after-tax income enters and neither hours worked nor productivity appears separately.

What remains to be determined is individual and aggregate effective labor supply. Without further loss of generality, I assume $G(h_{it}, n_{it}) = h_{it}\frac{n_{it}^{1+\gamma}}{1+\gamma}$. This functional form simplifies the household problem in the stationary equilibrium as h_{it} drops out from the first-order condition and all households supply the same number of hours $n_{it} = N(w_t)$. Total effective labor input, $\int n_{it}h_{it}di$, is hence also equal to $N(w_t)$ because I normalized $\int h_{it}di = 1$.[100]

Households also receive profit income from union profits Π_t^U or firms profits Π_t^{fi} as workers or entrepreneurs, respectively. Both profits get taxed at rate τ_t. What is more, households may receive *non-distortionary* targeted transfer as minimum income benefits tr_{it} as well as lump-sum transfers, Tr_t. All together, after-tax non-capital income, plugging in the optimal supply of hours, is then:

$$y_{it} = \left[(1-\tau_t)w_t/p_t^c(a_{it}^C)\right]^{\frac{1+\gamma}{\gamma}} h_{it} + \mathbb{I}_{h_{it}\neq 0}(1-\tau_t)\Pi_t^U + \mathbb{I}_{h_{it}=0}(1-\tau_t)\Pi_t^{fi} + tr_{it} + Tr_t. \tag{A.8}$$

Consumption, savings, and portfolio choice

Given this labor income, households optimize inter-temporally subject to their budget constraint expressed in terms of physical consumption goods:

$$p_t^c(a_{it}^C)c_{it} + b_{it+1} + q_t k_{it+1} = y_{it} + b_{it}\frac{R(b_{it}, R_t^b)}{\pi_t^{core}} + (q_t + r_t)k_{it}, \quad k_{it+1} \geq 0,\, b_{it+1} \geq \underline{B} \tag{A.9}$$

b_{it} is real bond holdings, k_{it} is the amount of illiquid assets, q_t is the price of these assets, r_t is their dividend, $\pi_t^{core} = \frac{P_t}{P_{t-1}}$ is realized average core inflation (inflation of physical goods,

[100]This means that I can read off average productivity risk from the estimated income risk series in the literature. Without scaling the labor dis-utility by productivity, I would need to translate productivity risk to income risk through the endogenous hour response.

i.e., without carbon dioxide), and R is the gross nominal interest rate on bonds, which depends on the portfolio position of the household and the central bank's interest rate R_t^b, which is set one period before.

All households that do not participate in the capital market ($k_{it+1} = k_{it}$) still obtain dividends and can adjust their bond holdings. Depreciated capital has to be replaced for maintenance, such that the dividend, r_t, is the net return on capital. Holdings of bonds have to be above an exogenous debt limit $\underline{B}$, and holdings of capital have to be non-negative.

Substituting the expression $c_{it} = x_{it} + \frac{1}{1+\gamma}\left[(1-\tau_t)w_t/p_t^c(a_{it}^C)\right]^{\frac{1+\gamma}{\gamma}} h_{it}$ for consumption, I obtain the budget constraint for the composite leisure-consumption good:

$$p_t^c(a_{it}^C)x_{it} + b_{it+1} + q_t k_{it+1} = b_{it}\frac{R(b_{it}, R_t^b)}{\pi_t^{core}} + (q_t + r_t)k_{it} + z_{it}, \quad k_{it+1} \geq 0,\, b_{it+1} \geq \underline{B}, \tag{A.10}$$

where $z_{it} = \frac{\gamma}{1+\gamma}\left[(1-\tau_t)w_t/p_t^c(a_{it}^C)\right]^{\frac{1+\gamma}{\gamma}} h_{it} + \mathbb{I}_{h_{it}\neq 0}(1-\tau_t)\Pi_t^U + \mathbb{I}_{h_{it}=0}(1-\tau_t)\Pi_t^{fi} + tr_{it} + Tr_t$ is income corrected for the dis-utility of labor.

Households make their savings choices and their portfolio choice between liquid bonds and illiquid capital in light of a capital market friction that renders capital illiquid because participation in the capital market is random and i.i.d. in the sense that only a fraction, λ, of households are selected to be able to adjust their capital holdings in a given period. This means that I specify:

$$R(b_{it}, R_t^b) = \begin{cases} R_t^b & \text{if } b_{it} \geq 0 \\ R_t^b + \bar{R} & \text{if } b_{it} < 0 \end{cases}. \tag{A.11}$$

The extra wedge for unsecured borrowing, $\bar{R}$, creates a mass of households with zero unsecured credit but with the possibility to borrow, though at a penalty rate.

Since a household's saving decision—(b'_a, k') for the case of adjustment and (b'_n, k') for non-adjustment—will be some non-linear function of that household's wealth and productivity, inflation and all other prices will be functions of the domestic joint distribution, Θ_t, of (b, k, h) in t and the foreign joint distribution, Θ_t^*. This makes Θ and Θ^* state variables of the household's planning problem and these distributions evolve as a result of the economy's reaction to aggregate shocks. For simplicity, I summarize all effects of aggregate state variables, including the distributions of wealth and income, by writing the dynamic planning problem with time-dependent continuation values.

This leaves me with three functions that characterize the household's problem: value function V^a for the case where the household adjusts its capital holdings, the function V^n

for the case in which it does not adjust, and the expected continuation value, $\mathbb{W}$, over both:

$$\begin{aligned} V_t^a(b,k,h,a^C) &= \max_{k',b'_a} u[x(b,b'_a,k,k',h,a^C)] + \beta\mathbb{E}_t\mathbb{W}_{t+1}(b'_a,k',h,a^C) \\ V_t^n(b,k,h,a^C) &= \max_{b'_n} u[x(b,b'_n,k,k,h,a^C)] + \beta\mathbb{E}_t\mathbb{W}_{t+1}(b'_n,k,h,a^C) \\ \mathbb{W}_{t+1}(b',k',h,a^C) &= \lambda V_{t+1}^a(b',k',h,a^C) + (1-\lambda)V_{t+1}^n(b',k,h,a^C). \end{aligned} \tag{A.12}$$

Expectations about the continuation value are taken with respect to all stochastic processes conditional on the current states, i.e., over both human capital, h, and carbon intensity, a^C. Maximization is subject to the corresponding budget constraint.

A.1.2 Firm sector

The firm sector consists of four sub-sectors: (a) a labor sector composed of unions that differentiate raw labor and labor packers who buy differentiated labor and then sell labor services to intermediate goods producers, (b) intermediate goods producers who hire labor services and rent out capital and buy energy to produce goods, (c) final goods producers who differentiate intermediate goods and then sell them to households and to (d) capital goods producers, who turn bundled goods into capital goods.

When profit maximization decisions in the firm sector require inter-temporal decisions (i.e. in price and wage setting and in producing capital goods), I assume for tractability that they are delegated to a mass-zero group of households (managers) that are risk-neutral and compensated by a share in profits. They do not participate in any asset market and have the same discount factor as all other households. Since managers are a mass-zero group in the economy, their consumption does not show up in any resource constraint, and all but the unions' profits go to the entrepreneur households (whose $h = 0$). Union profits go lump-sum to worker households.

Labor packers and unions

Worker households sell their labor services to a mass-n_A continuum of unions indexed by j, each of whom offers a different variety of labor to labor packers who then provide labor services to intermediate goods producers. Labor packers produce final labor services according to the production function

$$N_t = \left(\int_0^{n_A} \hat{n}_{jt}^{\frac{\eta_W - 1}{\eta_W}} dj\right)^{\frac{\eta_W}{\eta_W - 1}}. \tag{A.13}$$

out of labor varieties $\hat{n}_{jt}$. Cost minimization by labor packers implies that each variety of labor, each union j, faces a downward-sloping demand curve

$$\hat{n}_{jt} = \left(\frac{W_{jt}}{W_t^{fi}}\right)^{-\eta_W} N_t \tag{A.14}$$

where W_{jt} is the nominal wage set by union j and W_t^{fi} is the nominal wage at which labor packers sell labor services to final goods producers. Since unions have market power, they pay the households a wage lower than the price at which they sell labor to labor packers. Given the nominal wage W_t at which they buy labor from households and given the nominal wage index W_t^{fi}, unions seek to maximize their discounted stream of profits. However, they face a Calvo (1983) type adjustment friction with indexation with the probability λ_w to keep wages constant. They therefore maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \lambda_w^t \frac{W_t^{fi}}{P_t} N_t \left\{ \left(\frac{W_{jt}(\bar{\pi}_W)^t}{W_t^{fi}} - \frac{W_t}{W_t^{fi}}\right) \left(\frac{W_{jt}(\bar{\pi}_W)^t}{W_t^{fi}}\right)^{-\eta_W} \right\}. \tag{A.15}$$

by setting W_{jt} in period t and keeping it constant except for indexation to π_W, the steady state wage inflation rate.

Since all unions are symmetric, I focus on a symmetric equilibrium and obtain the linearized wage Phillips curve from the corresponding first-order condition as follows, leaving out all terms irrelevant at a first-order approximation around the stationary equilibrium:

$$\log\left(\frac{\pi_t^W}{\bar{\pi}^W}\right) = \beta \mathbb{E}_t \log\left(\frac{\pi_{t+1}^W}{\bar{\pi}^W}\right) + \kappa_w \left(mc_t^w - \frac{1}{\mu^W}\right), \tag{A.16}$$

with $\pi_t^W := \frac{W_t^{fi}}{W_{t-1}^{fi}} = \frac{w_t^{fi}}{w_{t-1}^{fi}} \pi_t^{CPI}$ being domestic wage inflation, w_t and w_t^{fi} being the respective *real* wages for households and firms, $mc_t^w = \frac{w_t}{w_t^{fi}}$ is the mark-down of wages the unions pay to households, W_t, relative to the wages charged to firms, W_t^{fi} and $\kappa_w = \frac{(1-\lambda_w)(1-\lambda_w \beta)}{\lambda_w}$. Union profits paid to workers therefore are $\Pi_t^U = (w_t^{fi} - w_t)N_t$.

Final goods producers

Similar to unions, final goods producers differentiate the homogeneous intermediate goods and set prices. They buy the intermediate good at the nominal price, MC_t. As I do for unions, I assume price adjustment frictions à la Calvo (1983) with indexation.

Under this assumption, the firms' managers maximize the present value of real profits given this price adjustment friction, i.e., they maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \lambda_Y^t (1-\tau_t) \left(\frac{p_{jt}(\bar{\pi})^t}{P_t} - \frac{MC_t}{P_t}\right) Y_t^d(j) \tag{A.17}$$

with a time-constant discount factor.

The corresponding first-order condition for price setting implies a domestic Phillips curve

$$\log\left(\frac{\pi_t}{\bar{\pi}}\right) = \beta\mathbb{E}_t \log\left(\frac{\pi_{t+1}}{\bar{\pi}}\right) + \kappa_Y\left(mc_t - \frac{1}{\mu^Y}\right) \tag{A.18}$$

where I again dropped all terms irrelevant for a first-order approximation and have $\kappa_Y = \frac{(1-\lambda_Y)(1-\lambda_Y\beta)}{\lambda_Y}$. Here, $\pi_t := \frac{P_t}{P_{t-1}}$, is the gross producer price inflation rate, i.e., the gross inflation rate of the physical good, $mc_t := \frac{MC_t}{P_t}$ are the real marginal costs, $\bar{\pi}$ is steady-state inflation, and $\frac{1}{\mu^Y} = \frac{\eta-1}{\eta}$ is the target markup. Profits paid to entrepreneurs therefore are $\Pi_t^F = (1 - mc_t)Y_t$.

Intermediate goods producers

Intermediate goods are produced with a constant returns to scale production function:

$$Y_t = \left((1-a_P)^{\frac{1}{\sigma_P}}\ Y_t^{P\,\frac{\sigma_P-1}{\sigma_P}} + a_P{}^{\frac{1}{\sigma_P}}\ \left(E_t^Y\right)^{\frac{\sigma_P-1}{\sigma_P}}\right)^{\frac{\sigma_P}{\sigma_P-1}},\ \text{where } Y_t^P = (u_t K_t^s)^\alpha\ N_t^{1-\alpha}. \tag{A.19}$$

Production combines physical production Y_t^P using capital K_t with capacity utilization u_t, labor N_t, and carbon dioxide E_t^Y. The coefficient α is the capital share, the coefficient σ_P captures the (short-run) substitutability of carbon dioxide in the production process, and a_P is the carbon dioxide share of production in normal times. Using capital with an intensity higher than normal increases depreciation of capital according to $\delta(u_t) = \delta_0 + \delta_1(u_t - 1) + \delta_2/2(u_t - 1)^2$, which, assuming $\delta_1, \delta_2 > 0$, is an increasing and convex function of utilization. Without loss of generality, capital utilization in the steady state is normalized to 1, so that δ_0 denotes the steady-state depreciation rate of capital goods.

Let mc_t be the relative price at which the intermediate good is sold to final goods producers. The intermediate goods producer maximizes profits,

$$mc_t Y_t - w_t^{fi} N_t - [r_t^F + q_t\delta(u_t)]K_t - (p_t^E - \tau_t^E)E_t^Y, \tag{A.20}$$

where r_t^F and q_t are the rental rate of firms and the (producer) price of capital goods, respectively. The intermediate goods producer operates in perfectly competitive markets, such that the real wage and the user costs of capital are determined by the following equations:

$$MPK_t = mc_t\,(1-a_P)^{\left(\frac{1}{\sigma_P}\right)}\,\alpha\left(\frac{K_t}{N_t}\right)^{(\alpha-1)}\left(\frac{Y_t}{Y_t^P}\right)^{\left(\frac{1}{\sigma_P}\right)}, \tag{A.21}$$

$$r_t = 1 + MPK_t u_t - q_t\delta(u_t), \tag{A.22}$$

$$w_t^{fi} = mc_t \left(1 - a_P\right)^{\left(\frac{1}{\sigma_P}\right)} (1-\alpha) \left(\frac{u_t K_t}{N_t}\right)^{\alpha} \left(\frac{Y_t}{Y_t^p}\right)^{\left(\frac{1}{\sigma_P}\right)}, \tag{A.23}$$

$$p_t^E - \tau_t^E = mc_t a_P^{\left(\frac{1}{\sigma_P}\right)} \left(\frac{Y_t}{E_t^Y}\right)^{\left(\frac{1}{\sigma_P}\right)}. \tag{A.24}$$

Here MPK is the marginal product of capital services. I assume that utilization is decided by the owners of the capital goods, taking the aggregate supply of capital services as given. The optimality condition for utilization is given by

$$MPK_t = q_t[\delta_1 + \delta_2(u_t - 1)] \tag{A.25}$$

i.e., capital owners increase utilization until the marginal maintenance costs equal the marginal product of capital services.

Capital goods producers

Capital goods producers transform the physical good, investment I_t, into capital. They take the relative price of capital goods, q_t, as given in deciding about their output, i.e., they maximize[101]

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t I_t \left\{ q_t \left[1 - \frac{\phi}{2}\left(\log \frac{I_t}{I_{t-1}}\right)^2\right] - 1 \right\}. \tag{A.26}$$

Optimality of the capital goods production requires (again dropping all terms irrelevant up to first order)

$$q_t \left[1 - \phi \log \frac{I_t}{I_{t-1}}\right] = 1 - \beta \mathbb{E}_t \left[q_{t+1} \psi \log\left(\frac{I_{t+1}}{I_t}\right)\right], \tag{A.27}$$

and each capital goods producer will adjust its production until (C.50) is fulfilled.

Since all capital goods producers are symmetric, I obtain the law for motion for aggregate capital as

$$K_t - (1 - \delta(u_t))K_{t-1} = \left[1 - \frac{\phi}{2}\left(\log \frac{I_t}{I_{t-1}}\right)^2\right] I_t \tag{A.28}$$

The functional form assumption implies that investment adjustment costs are minimized and equal to 0 in the steady state.

[101] As I use a first order approximation changes in the stochastic discount factor are irrelevant. So are changes in the relative price $p_t(a^C)$ of the physical to the final consumption good.

A.1.3 Government Sector

There is a monetary authority and a fiscal authority. The monetary authority controls the nominal interest rate on liquid assets, while the fiscal authorities issue government bonds to finance deficits, choose the average tax rate, make expenditures for government consumption and their transfer system, and receive revenue from selling carbon certificates.

Monetary Union

I assume that monetary policy sets the nominal interest rate following a Taylor (1993)-type rule with interest rate smoothing:

$$\frac{R^b_{t+1}}{\bar{R}^b} = \left(\frac{R^b_t}{\bar{R}^b}\right)^{\rho_R} \left(\frac{\pi_t}{\bar{\pi}}\right)^{(1-\rho_R)\theta_\pi} \left(\frac{Y_t}{Y_{t-1}}\right)^{(1-\rho_R)\theta_Y} . \tag{A.29}$$

The coefficient $\bar{R}^b \geq 0$ determines the nominal interest rate in the steady state. The coefficients $\theta_\pi, \theta_Y \geq 0$ govern the extent to which the central bank attempts to stabilize producer price inflation and output growth. $\rho_R \geq 0$ captures interest rate smoothing.

Fiscal Policy

The budget constraint of the fiscal policy reads

$$G_t + TR_t = B_{t+1} + T_t - \frac{R^b_t}{\pi^{CPI}_t} B_t + T^E_t. \tag{A.30}$$

Hence, the government has expenditure for government spending, G_t, aggregate spending on its transfer system specified below, TR_t, and repaying its debt, B_t. It finances its expenditures by issuing new debt, collecting tax revenue, T_t, and by collecting the revenue from selling carbon certificates, $T^E_t = p^E_t * E_t$. Tax revenue is

$$T_t = \tau_t(w_t N_t + \mathbb{I}_{h_{it}=0}\Pi^{fi}_t + \mathbb{I}_{h_{it}\neq 0}\Pi^U_t). \tag{A.31}$$

I assume that the average tax rate is a feedback function of government debt:

$$\frac{\tau_t}{\bar{\tau}} = \left(\frac{\tau_{t-1}}{\bar{\tau}}\right)^{\rho_\tau} \left(\frac{B_{t+1}}{\bar{B}}\right)^{(1-\rho_\tau)\gamma^\tau_B} . \tag{A.32}$$

where γ^τ_B governs the speed with which debt returns to its target.

Targeted Transfer System

The targeted transfer system follows the design in Bayer et al. (2023). It provides additional resources if net labor income $w_t n_t h_{it}$ falls short of some target level. For simplicity, I assume that these transfers are non-distortionary for the labor supply decision. In particular, I assume that transfers are paid to households according to the following scheme:

$$tr_{it} = \max\{0, a_1\bar{y} - a_2(1-\tau_t)w_t h_{it} n_{it}\}, \tag{A.33}$$

where $\bar{y}$ is the median income and $0 \leq a_1, a_2 \leq 1$. Thus, transfers decrease in individual income with a transfer withdrawal rate of a_2 and no transfers are paid to households whose net labor income $(1-\tau_t)w_t h_{it} n_{it} \geq \frac{a_1}{a_2}\bar{y}$. Total transfer payments are then

$$TR_t = \mathbb{E}_t tr_{it} + Tr_t, \tag{A.34}$$

where again, the expectation operator is the cross-sectional average.

A.1.4 Carbon dioxide, goods, bonds, capital, and labor market clearing

The market for carbon certificates clears, when total carbon dioxide emission, consisting of household and firm carbon dioxide emission, equals the exogenous supply of carbon certificates:

$$E_t = E_t^C + E_t^Y. \tag{A.35}$$

The labor market clears at the competitive wage given in (C.46). The bond markets clear whenever the following equation holds:

$$B_{t+1} = B^d(p_t^E, Tr_t, R_t^b, r_t, q_t, \Pi_t^{fi}, \Pi_t^U, w_t, \pi_t, \tau_t, \Theta_t, \Theta_t^*, \mathbb{W}_{t+1}) \quad := \mathbb{E}_t[\lambda \mathbb{B}_{a,t} + (1-\lambda)\mathbb{B}_{n,t}], \tag{A.36}$$

where $\mathbb{B}_{a,t}$, $\mathbb{B}_{n,t}$ are functions of the states (b, k, h, a^c), and depend on how the households value asset holdings in the future, $\mathbb{W}_{t+1}$, and the current set of prices (and tax rates) $(p_t^E, Tr_t, R_t^b, r_t, q_t, \Pi_t^{fi}, \Pi_t^U, w_t, \pi_t^{CPI}, \tau_t)$. Future prices do not show up because I can express the value functions such that they summarize all relevant information on the expected future price paths. Expectations in the right-hand-side expression are taken w.r.t. the distributions $\Theta_t(b, k, h, a^c)$. Equilibrium requires the total *net* amount of bonds the household sectors demand to equal the supply of government bonds. In gross terms, there are more liquid assets in circulation as some households borrow up to $\underline{B}$.

In addition, the market for capital has to clear:

$$\begin{aligned} K_{t+1} &= K^d(p_t^E, Tr_t, R_t^b, r_t, q_t, \Pi_t^{fi}, \Pi_t^U, w_t, \pi_t^{CPI}, \tau_t, \Theta_t, \Theta_t^*, \mathbb{W}_{t+1}) \\ &:= \mathbb{E}_t[\lambda(\mathbb{K}_t) + (1-\lambda)(k)] \end{aligned} \tag{A.37}$$

where the first equation stems from competition in the production of capital goods, and the second equation defines the aggregate supply of funds from households - both those that trade capital, $\lambda(\mathbb{K}_t)$ and those that do not, $(1-\lambda)(k)$. Again $\mathbb{K}_t$ is a function of the current prices and continuation values.

Finally, goods market clearing requires:

$$Y_t = C_t + I_t + BD_t\bar{R} + G_t. \tag{A.38}$$

A.1.5 Equilibrium

A sequential equilibrium with recursive planning in my Climate-HANK model is a sequence of policy functions $\{\mathbb{X}_{at}, \mathbb{X}_{nt}, \mathbb{B}_{at}, \mathbb{B}_{nt}, \mathbb{K}_t\}$, a sequence of value functions $\{V_t^a, V_t^n\}$, a sequence of prices $\{p_t^E, \tau_t^E, Tr_t, w_t, w_t^{fi}, \Pi_t^U, \Pi_t^{fi}, q_t, r_t, R_t^b, \pi_t^{CPI}, \pi_t^W, \tau_t\}$, a sequence of carbon certificates, $\{E_t\}$, aggregate capital, labor supply, distributions Θ_t over individual asset holdings and productivity, and expectations for the distribution of future prices, Γ, such that

1. Given the functionals $\mathbb{E}_t\mathbb{W}_{t+1}$ for the continuation value and period-t prices, policy functions $\{\mathbb{X}_{at}, \mathbb{X}_{nt}, \mathbb{B}_{at}, \mathbb{B}_{nt}, \mathbb{K}_t\}$ solve the households' planning problem; and given the policy functions $\{\mathbb{X}_{at}, \mathbb{X}_{nt}, \mathbb{B}_{at}, \mathbb{B}_{nt}, \mathbb{K}_t\}$ and prices, the value functions $\{V_t^a, V_t^n\}$ are a solution to the Bellman equation.

2. Distributions of wealth and income evolve according to households' policy functions.

3. All markets clear in every period, interest rates on bonds are set according to the central bank's Taylor rule, fiscal policies are set according to the fiscal rules, and stochastic processes evolve according to their law of motion.

4. Expectations are model consistent.

I solve the model by using the perturbation method in Bayer et al. (2020b).

A.2 Calibration

I calibrate the economy to match German data. To this end, I match the wealth distributions. Table C.1 shows the calibration choices required for our calibration strategy which is described in 1.3. The rest of the parameters are calibrated by matching long-run averages and using standard parameters from the literature. Table C.2 summarizes my calibration of those parameters. I calibrate to quarterly frequency.

The labor share in production, $(1-\alpha)$, is 68% corresponding to a labor income share of 62%, given a markup of 10% due to an elasticity of substitution between differentiated goods of 11. The elasticity of substitution between labor varieties is also set to 11, yielding a wage markup of 10%. The parameter δ_1 that governs the cyclicality of utilization is set to 5.0. The investment adjustment cost parameter is set to 4.0. I set the Calvo parameters for price and wage adjustment probability both to 0.25. All these parameter choices are standard values in the literature.

Table A.1: Calibration—Asymmetric Parameters

	Description	**Germany**	**Source/Target**
a_1	Transfer level	0.5	German MIB system
a_2	Transfer withdrawal rate	0.8	German MIB system
G/Y	Gov. cons. share	0.20	German data
σ_h	STD labor inc.	0.135	German income data
β	Discount factor	0.9823	Six wealth targets
λ	Portfolio adj. prob.	0.071	Six wealth targets
ζ	Trans. prob. from W to E	0.001	Six wealth targets
ι	Trans prob. E to W	0.0625	Six wealth targets
$\bar{R}$	Borrowing penalty	0.029	Six wealth targets
B_{min}/Y	Borrowing limit	1.7	Six wealth targets

I set relative risk aversion, ξ, to 4, following Kaplan and Violante (2014) and the Frisch elasticity, γ to 0.5 following Chetty et al. (2011). The persistence of idiosyncratic income shocks is set to $\rho_h = 0.9815$. The stationary equilibrium real rate(-growth difference) is set to a net rate of zero.

The steady-state tax level is set to 0.3. I assume that monetary policy only targets inflation, as this is the primary mandate of the ECB, and set the Taylor coefficient to 1.25 and the smoothing parameter to 0.85. The steady-state inflation is zero.

A.3 Further results

Table A.2: Rest of Calibration

	Description	Value	Source/Target
Firms			
$1-\alpha$	Share of labor	0.68	62% lab. income
η	Elast. of substitution	11	10% Price markup
η_W	Elast. of substitution	11	10% Wage markup
κ	Price adj. prob.	0.25	1 year avg. price duration
κ_W	Wage adj. prob.	0.25	1 year avg. wage duration
ϕ	Inv. adj. cost	4.0	Bayer et al. (2020b)
δ_0	Depreciation rate	0.018	Bayer et al. (2020b)
δ_1	Depr. rate increase	5.0	Bayer et al. (2020b)
Households			
ξ	Risk aversion	4	Kaplan and Violante (2014)
γ	Inv. Frisch elast.	2	Chetty et al. (2011)
Government			
$\bar{\tau}$	Tax rate	0.3	Standard value
ρ_R	Pers. in Tax rule	0.9	standard value
γ_B^τ	Reaction to debt.	0.85	standard value
$\bar{R}^b$	Gross interest rate	1.00	zero interest-growth difference
ρ_R	Pers. in Taylor rule	0.85	standard value
θ_π	Reaction to Infl.	1.25	standard value
θ_Y	Reaction to Output	0	ECB mandate

Table A.3: Calibrated Model v Data

			Model		Data	
			F	**H**	**ITA**	**GER**
Steady state	Assets	Debt (% of output)	132	71	132	71
		Capital-Output-Ratio	3.3	3.2	3.3	3.2
(targeted)	Distribution	Wealth gini	0.60	0.72	0.61	0.73
		Top-10% wealth share	0.43	0.55	0.44	0.52
		Bottom-50% wealth share	0.10	0.01	0.09	0.02
		Borrowers	0.08	0.18	0.08	0.18

Notes: Model predictions based on baseline calibration, see Appendix A.2 for details. Microdata based on the 2017 wave of the Household Finance and Consumption survey of the ECB. Macro data from Eurostat. Quantities are measured in real per capita terms, yoy changes; sample: 1999Q1-2022Q2.

Figure A.1: Response to carbon shock, transfer repayment

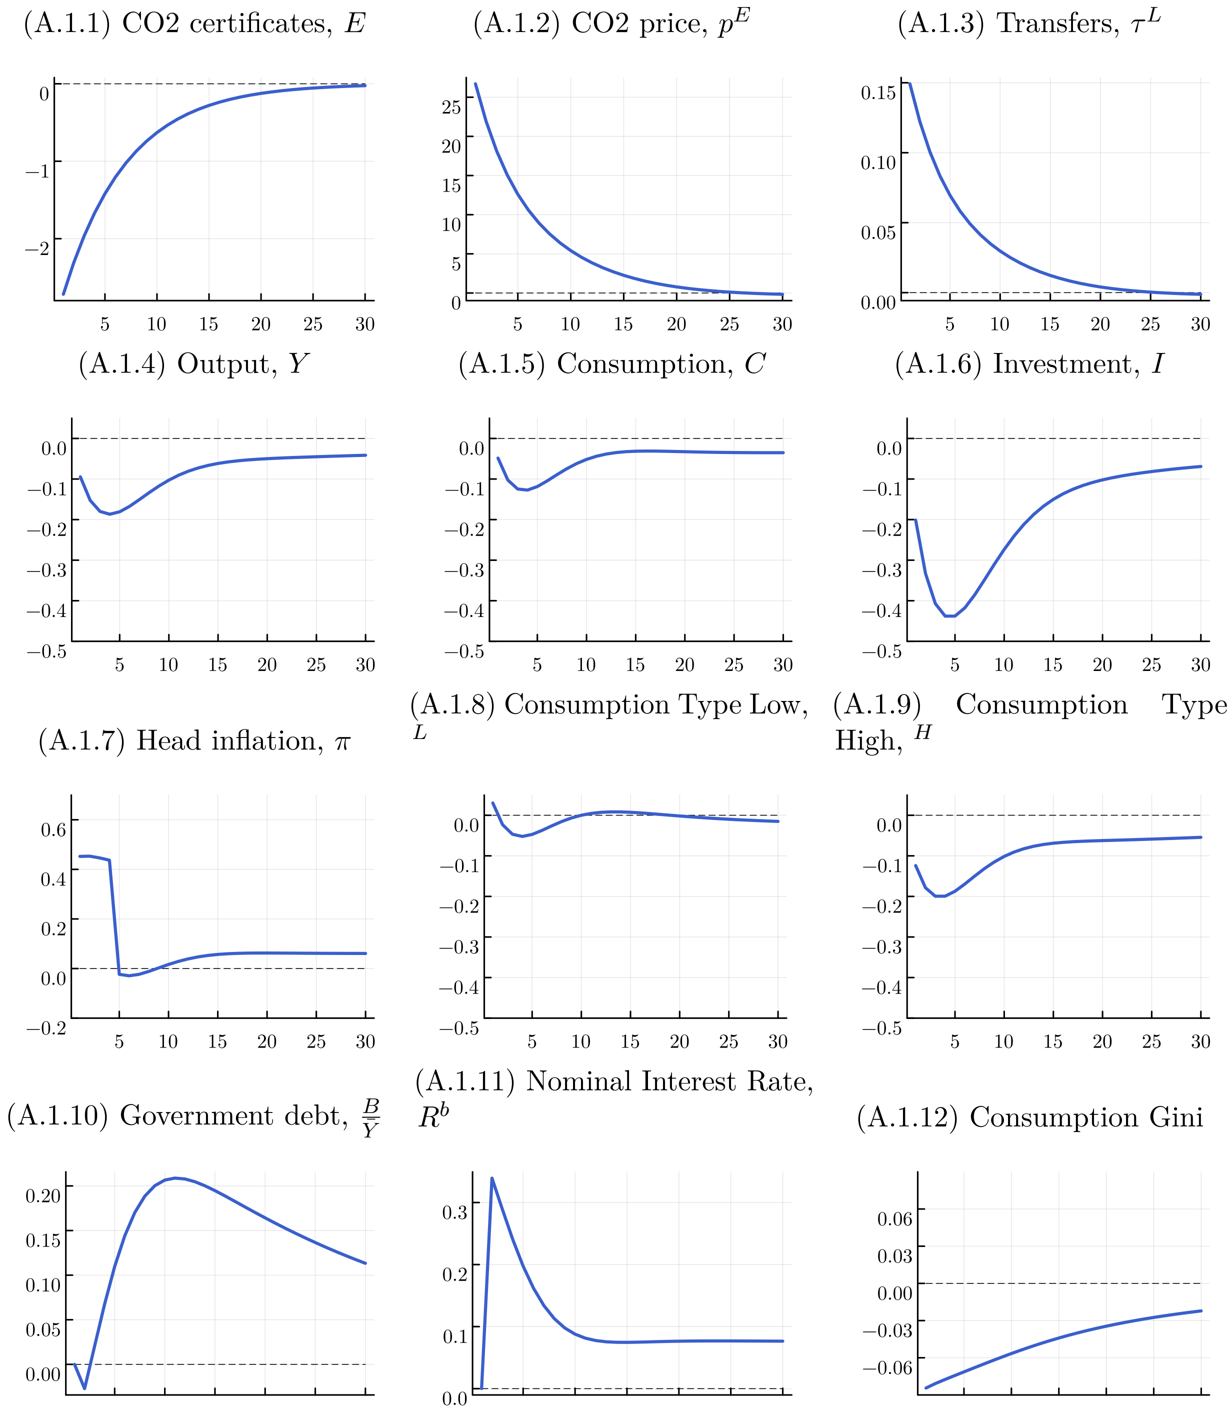

Impulse responses after an adverse carbon shock. Y-axis:Percentage deviation from steady state, percentage points in case of inflation and interest rate, and tax level in percentage in case of taxes. X-axis: Quarters.

Figure A.2: Response to carbon shock, tax repayment

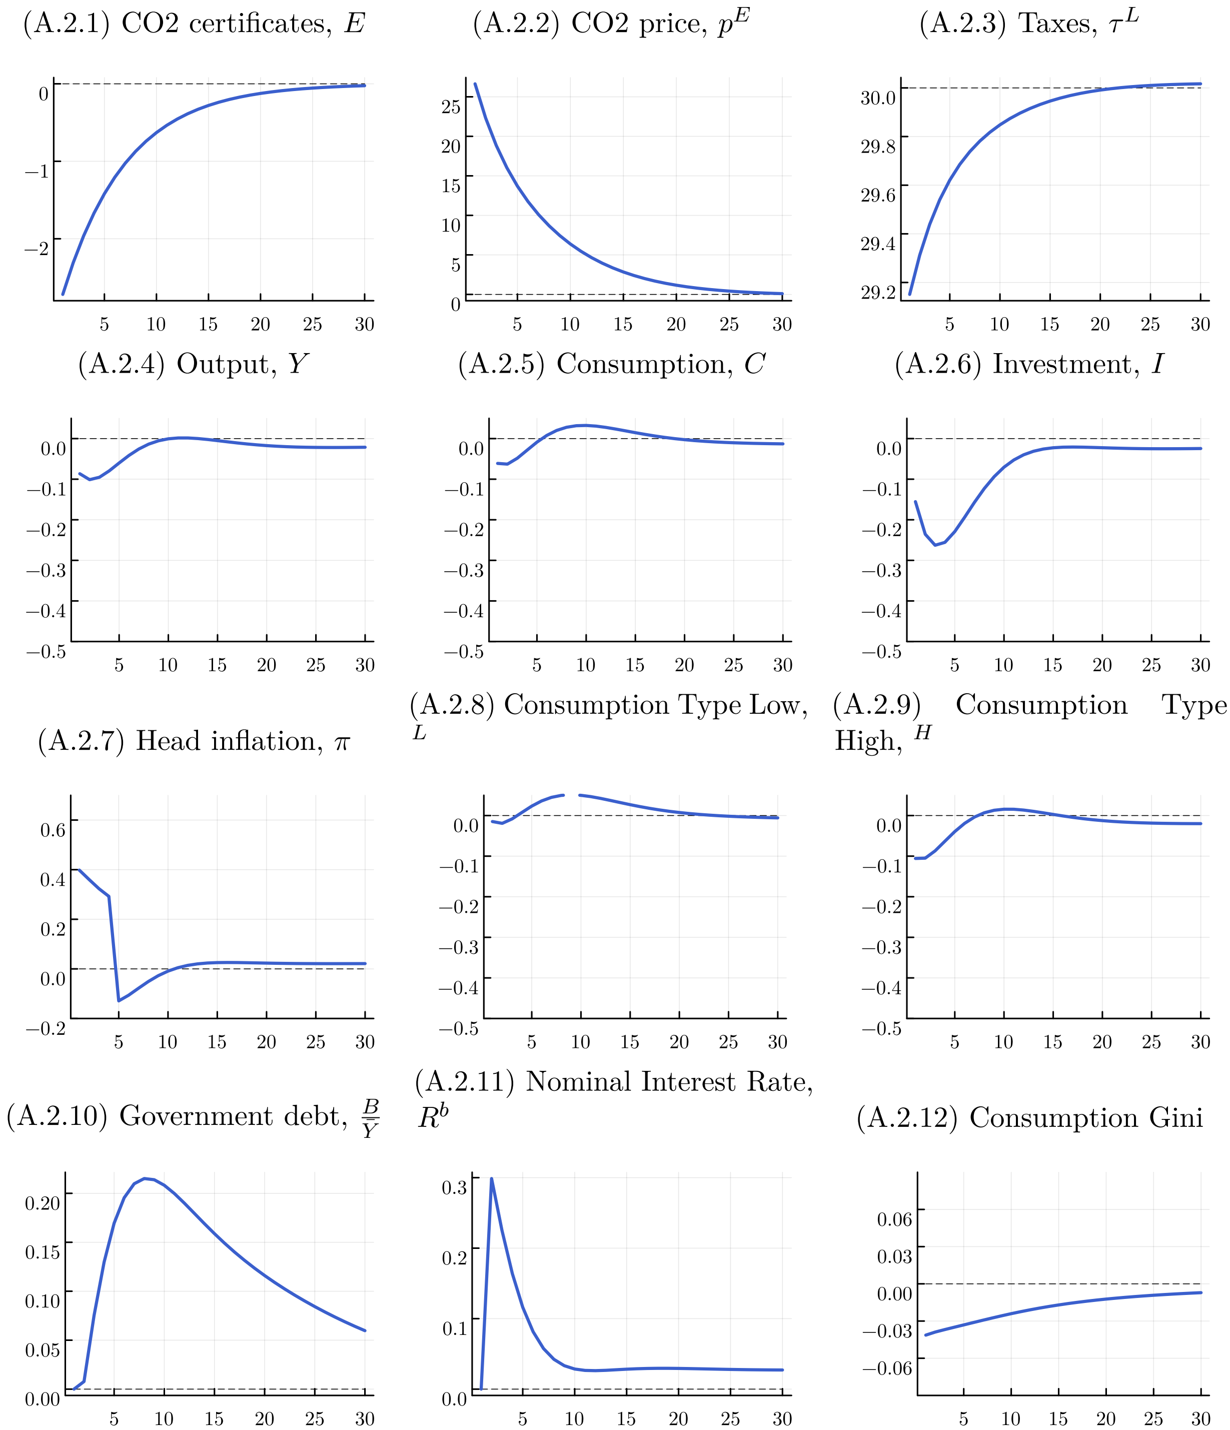

(A.2.1) CO2 certificates, E
(A.2.2) CO2 price, p^E
(A.2.3) Taxes, τ^L
(A.2.4) Output, Y
(A.2.5) Consumption, C
(A.2.6) Investment, I
(A.2.7) Head inflation, π
(A.2.8) Consumption Type Low, ^{L}
(A.2.9) Consumption Type High, ^{H}
(A.2.10) Government debt, $\frac{B}{Y}$
(A.2.11) Nominal Interest Rate, R^b
(A.2.12) Consumption Gini

Impulse responses after an adverse carbon shock. Y-axis:Percentage deviation from steady state, percentage points in case of inflation and interest rate, and tax level in percentage in case of taxes. X-axis: Quarters.

Figure A.3: Response to carbon shock, capital subsidy repayment

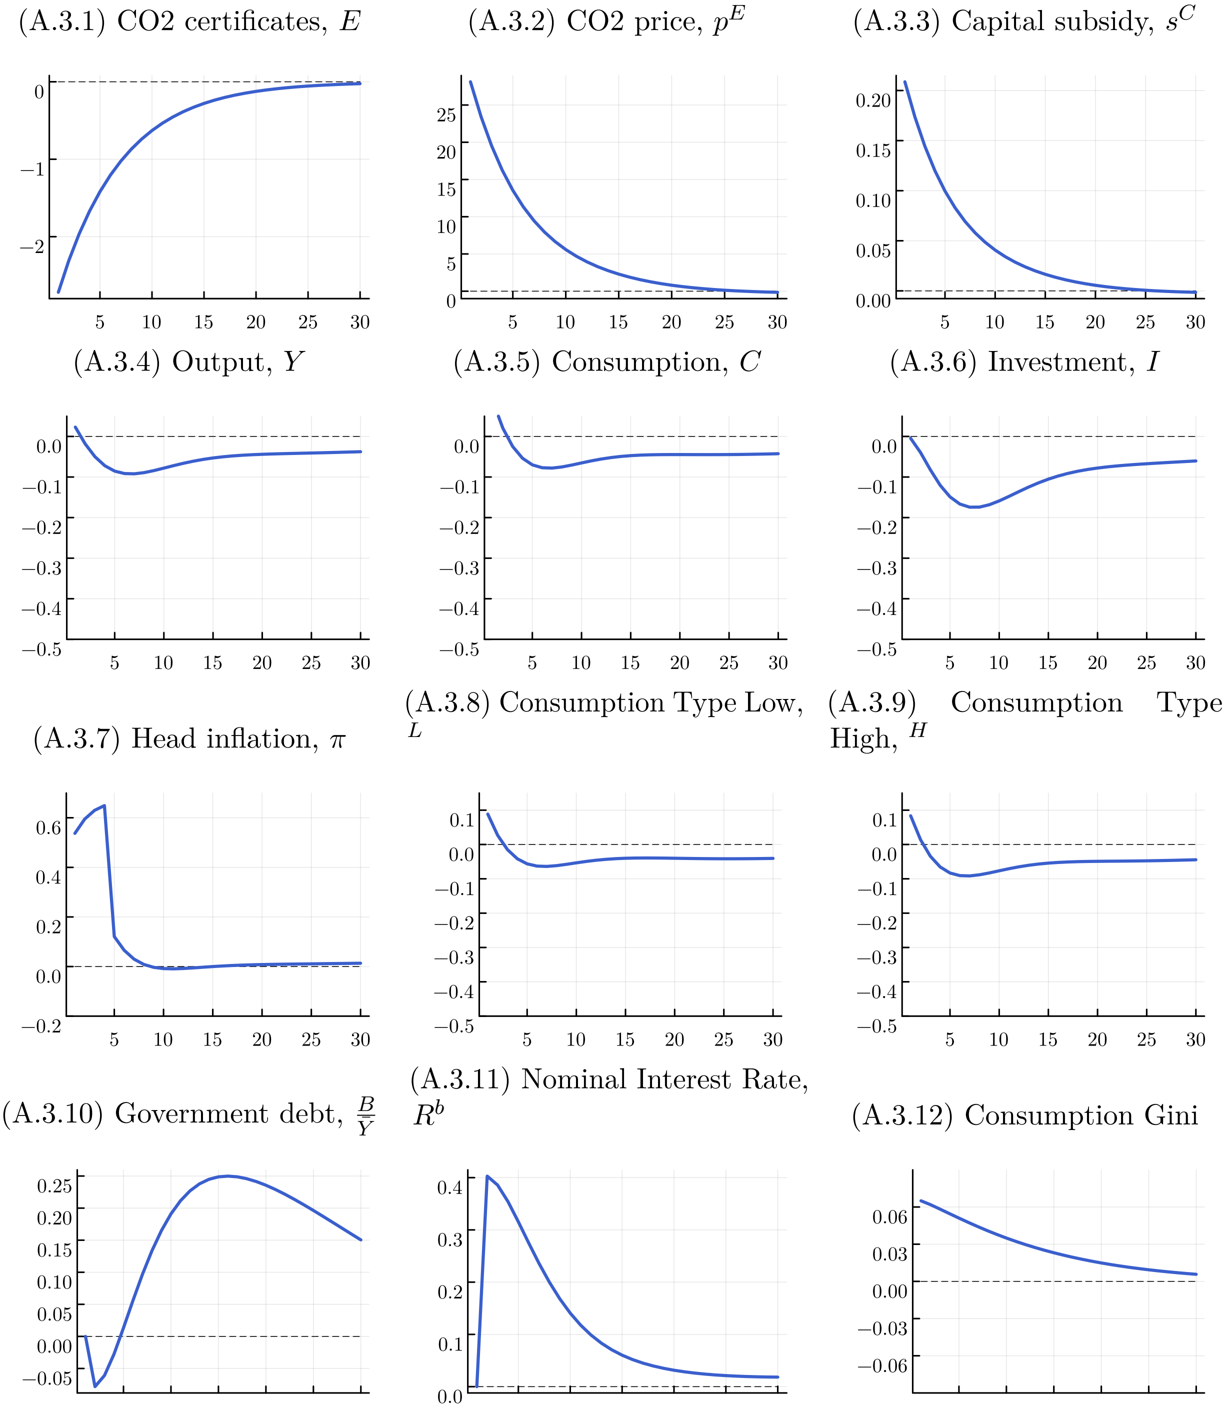

Impulse responses after an adverse carbon shock. Y-axis:Percentage deviation from steady state, percentage points in case of inflation, interest rate, and capital subsides. X-axis: Quarters.

Appendix B

Appendix for Chapter 2

B.1 Proof of Proposition 1

In this section, we prove Proposition 1 which states that HANK-UFP yields the same allocation as the one induced by the real interest rate path in the monetary policy experiment. Let us assume that the equilibrium path induced by the monetary policy experiment in Section (2.3) is

$$X^{MP} = \left\{ B_t^{d*}, C_t^*, L_t^*, Y_t^*, D_t^*, w_t^*, \pi_t^*, \tilde{p_t}^*/P_t^*, r_t^{MP}, \bar{\tau}^C, \bar{\tau}^L, Tr_t^{MP}, \bar{B} \right\}_{t=0}^{\infty},$$

with the individual behavior of each household given by

$$x_h^{MP} = \{b_{h,t+1}^*, c_{h,t}^*, l_{h,t}^*\}_{t=0}^{\infty}.$$

We now show that

$$X^{UFP} = \left\{ \frac{1+\tau_{t-1}^{C,UFP}}{1+\bar{\tau}^C} B_t^{d*}, C_t^*, L_t^*, Y_t^*, D_t^*, w_t^*, \pi_t^*, \tilde{p_t}^*/P_t^*, \bar{r}, \tau_t^{C,UFP}, \tau_t^{L,UFP}, Tr_t^{UFP}, B_t^{UFP} \right\}_{t=0}^{\infty}$$

$$\text{with } x_h^{UFP} = \left\{ \frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C} b_{h,t+1}^*, c_{h,t}^*, l_{h,t}^* \right\}_{t=0}^{\infty}$$

is an equilibrium path if $\tau_t^{L,UFP}, \tau_t^{C,UFP}$, and B_t^{UFP} satisfy conditions (2.20), (2.21), and (2.22), respectively.

Since neither the real interest rate nor the fiscal policy variables show up in any equilibrium condition of the firm side, it is sufficient to show that X^{UFP} satisfies the sequences of Euler equations (2.3), labor-leisure equations (2.4), and budget constraints (2.2) of each household as well as the government budget constraint. Without loss of generality, fix a household j. We now prove that if her behavior in the monetary policy experiment, $x_j^{MP} = \{b_{j,t+1}^*, c_{j,t}^*, l_{j,t}^*\}_{t=0}^{\infty}$, satisfies her sequences of Euler equations, labor leisure equations, and budget constraints in the monetary policy experiment, $x_j^{UFP} = \{\frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C} b_{j,t+1}^*, c_{j,t}^*, l_{j,t}^*\}_{t=0}^{\infty}$

does so in the HANK-UFP case.

Satisfying household's first-order conditions. Any path of consumption, $\{c^*_{j,t}\}_{t=0}^{\infty}$, that satisfies the sequence of Euler equations of household j with $\{r_t^{MP}\}_{t=0}^{\infty}$ and steady state tax rates, also satisfies the sequence of Euler equations of household j with interest rates in steady state and $\{\tau_t^{C,UFP}\}_{t=0}^{\infty}$ which satisfies condition (2.20). In addition, any paths of consumption and labor, $\{c^*_{j,t}, l^*_{j,t}\}_{t=0}^{\infty}$, that satisfy the sequence of labor-leisure equations of household j with steady state taxes, satisfy the sequence of labor-leisure equations of household j if $\{\tau_t^{L,UFP}, \tau_t^{C,UFP}\}_{t=0}^{\infty}$ satisfy condition (2.21).

Satisfying household's budget constraint. Next, we show that if x_j^{MP} satisfies the sequence of budget constraints of household j in the monetary policy experiment, x_j^{UFP} does so with HANK-UFP. To this end, it is convenient to look at her three income components separately:

$$c_{j,t} = \underbrace{\frac{1-\tau_t^L}{1+\tau_t^C} w_t l_{j,t} z_{j,t}}_{I} + \underbrace{\frac{D_t + Tr_t}{1+\tau_t^C}}_{II} + \underbrace{(1+r_t)\frac{b_{j,t}}{1+\tau_t^C} - \frac{b_{j,t+1}}{1+\tau_t^C}}_{III}. \tag{B.1}$$

We now show that each of these components is exactly the same with x_j^{MP} and X^{MP} as well as with x_j^{UFP} and X^{UFP}.

The labor income, term (I) in equation (B.1), is the same in the monetary policy experiment and in the HANK-UFP case, iff:

$$\frac{1-\bar{\tau}^L}{1+\bar{\tau}^C} w_t^* l_{j,t}^* z_{j,t} = \frac{1-\tau_t^{L,UFP}}{1-\tau_t^{C,UFP}} w_t^* l_{j,t}^* z_{j,t} \tag{B.2}$$

which holds, given that taxes are set consistent with condition (2.21).

The lump-sum income, term (II) in equation (B.1), is identical in the monetary policy experiment and in the HANK-UFP case, iff:

$$\frac{D_t^* + Tr_t^{MP}}{1+\bar{\tau}^C} = \frac{D_t^* + Tr_t^{UFP}}{1+\tau_t^{C,UFP}} \tag{B.3}$$

which holds if Tr_t^{UFP} is set according to condition (2.23). We will show below that this is indeed the transfer path arising in the HANK-UFP case.

Finally, the asset income, term (III) in equation (B.1), is the same iff:

$$(1+r_t^{MP})\frac{b^*_{j,t}}{1+\bar{\tau}^C} - \frac{b^*_{j,t+1}}{1+\bar{\tau}^C} = (1+\bar{r})\frac{b^*_{j,t}\frac{1+\tau_{t-1}^{C,UFP}}{1+\bar{\tau}^C}}{1+\tau_t^{C,UFP}} - \frac{b^*_{j,t+1}\frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C}}{1+\tau_t^{C,UFP}}. \tag{B.4}$$

Hence, we obtain equivalence iff:

$$(1+r_t^{MP})\frac{b_{j,t}^*}{1+\bar{\tau}^C} = (1+\bar{r})\frac{b_{j,t}^*\frac{1+\tau_{t-1}^{C,UFP}}{1+\bar{\tau}^C}}{1+\tau_t^{C,UFP}}$$

Using condition (2.20), we get

$$\begin{aligned}(1+r_t^{MP})\frac{b_{j,t}^*}{1+\bar{\tau}^C} &= (1+\bar{r})\frac{b_{j,t}^*\frac{(1+r_t^{MP})}{1+\bar{\tau}^C}}{1+\bar{r}} \\ &\iff b_{j,t}^* = b_{j,t}^*.\end{aligned}$$

Thus, equation (B.10) holds. Hence, x_j^{UFP} satisfies the sequence of budget constraints of household j with HANK-UFP if x_j^{MP} does so in the monetary policy experiment. Furthermore, this also implies that if the individual state of household j is $(b_{j,t}^*, z_{j,t})$ in a given t in the monetary policy experiment, it is $(\frac{1+\tau_{t-1}^C}{1+\bar{\tau}^C} b_{j,t}^*, z_{j,t})$ with HANK-UFP.

In sum, $x_{h,t}^{UFP}$ and X^{UFP} are consistent with each household's problem if $x_{h,t}^{MP}$ and X^{MP} are. That is, each household consumes and works the same with HANK-UFP and in the monetary policy experiment and saves the same amount in consumption value terms.

Satisfying government's budget constraint. We now show that if X^{MP} satisfies the government's budget constraint, X^{UFP} does so as well. In the monetary policy experiment, the government's budget constraint is given by:

$$Tr_t^{MP} + r_t^{MP}\bar{B} = T_t^{MP}. \tag{B.5}$$

In the HANK-UFP case, the government's budget constraint is given by:

$$Tr_t^{UFP} + (1+\bar{r})B_t^{UFP} = B_{t+1}^{UFP} + T_t^{UFP}.$$

We show that given Proposition 1, this is exactly the same as equation (B.5). Assuming the transfer path given by condition (2.23), plugging in condition (2.22), using condition (2.20) and rearranging yields:

$$Tr_t^{MP}\left(\frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C}\right) + D_t\left(\frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C} - 1\right) - T_t^{UFP} = \bar{B}\left(\frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C}\right)\left(-r_t^{MP}\right).$$

Multiplying by $\frac{1+\bar{\tau}^C}{1+\tau_t^{C,UFP}}$ yields

$$Tr_t^{MP} + D_t\left(1 - \frac{1+\bar{\tau}^C}{1+\tau_t^{C,UFP}}\right) - T_t^{UFP}\frac{1+\bar{\tau}^C}{1+\tau_t^{C,UFP}} = -r_t^{MP}\bar{B}.$$

This is the same as in the monetary policy experiment if

$$D_t\left(1-\frac{1+\bar{\tau}^C}{1+\tau_t^{C,UFP}}\right)-T_t^{UFP}\frac{1+\bar{\tau}^C}{1+\tau_t^{C,UFP}}=-T_t^{MP}. \tag{B.6}$$

Using the goods market clearing condition, $Y_t^*=C_t^*$, and the profit equation $D_t^*=Y_t^*-w_t^*L_t^*$, it now holds that (dropping the superscript UFP for the sake of readability):

$$\begin{aligned}
&D_t^*\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C}-1\right)*\frac{1+\tau_t^C}{1+\bar{\tau}^C}(\bar{\tau}^C C_t^*+\bar{\tau}^L w_t^* L_t^*)=\tau_t^C C_t^*+\tau_t^L w_t^* L_t^*\\
&\qquad\iff D_t^*(\tau_t^C-\bar{\tau}^C)+C_t^*(\bar{\tau}^C-\tau_t^C)=(\tau_t^L+\bar{\tau}^C\tau_t^L-\bar{\tau}^L-\tau_t^C\bar{\tau}^L)(C_t^*-d_t^*)\\
&\iff (\bar{\tau}^C-\tau_t^C-\tau_t^L-\bar{\tau}^C\tau_t^L+\bar{\tau}^L+\tau_t^C\bar{\tau}^L)C_t^*=(\bar{\tau}^C-\tau_t^C-\tau_t^L-\bar{\tau}^C\tau_t^L+\bar{\tau}^L+\tau_t^C\bar{\tau}^L)d_t^*
\end{aligned}$$

Thus, the government budget constraint is satisfied if:

$$\bar{\tau}^C-\tau_t^C-\tau_t^L-\bar{\tau}^C\tau_t^L+\bar{\tau}^L+\tau_t^C\bar{\tau}^L=0 \tag{B.7}$$

$$\begin{aligned}
&\iff \bar{\tau}^C-\tau_t^C-(1+\bar{\tau}^C)\tau_t^L+(1+\tau_t^C)\bar{\tau}^L=0\\
&\iff 1+\bar{\tau}^C-(1+\tau_t^C)+(1+\tau_t^C)\bar{\tau}^L=(1+\bar{\tau}^C)\tau_t^L\\
&\iff -(1-\bar{\tau}^L)(1+\tau_t^C)=(\tau_t^C-1)(1+\bar{\tau}^C)
\end{aligned}$$

$$\iff \frac{1-\tau_t^L}{1+\tau_t^C}=\frac{1-\bar{\tau}^L}{1+\bar{\tau}^C}. \tag{B.8}$$

Which holds given that τ_t^C and τ_t^L are set according to (2.21).

Consistency with optimal behavior of firms. Given the same households' behavior $\{c_{h,t}^*, l_{h,t}^*\}_{t=0}^{\infty}$ in both policy cases, the firms also face the same demand for goods and the same supply of labor. Hence, if $\{w_t^*, D_t^*, \pi_t^*, \tilde{p_t}^*/P_t^*\}_{t=0}^{\infty}$, are equilibrium paths in the monetary policy experiment, they are also equilibrium paths in the UFP case.

Market clearing conditions. From individual behavior x_h^{MP} and x_h^{UFP}, it follows that the sequence of distributions in the HANK-UFP case is $\left\{\Gamma_t^{UFP}(\frac{1+\tau_{t-1}^C}{1+\bar{\tau}^C}b,z)\right\}_{t=0}^{\infty}=\left\{\Gamma_t^{MP}(b,z)\right\}_{t=0}^{\infty}$. That is, if the asset position is adjusted by the consumption value, the distributions are equivalent. Hence, if the asset market clears in the monetary policy experiment, aggregate savings are $B_{t+1}^{d,UFP}=\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{B}$ with HANK-UFP which is equal to the supply of government bonds in the HANK-UFP case.

Given the same behavior of firms and the same consumption and labor supply of households, X^{UFP} also clears all other markets if X^{MP} clears all other markets. Thus, we have proven that UFP set according to Proposition 1 yields the same allocation as in the monetary policy experiment which implies that UFP and monetary policy are perfect substitutes in HANK.

B.2 Proof of Lemma 1

Policy-exposure to monetary policy. We derive each household's policy-exposure in the monetary policy experiment, $\Xi^{MP}_{h,t}$, which is defined as the net excess in resources for each household given that only policy variables change. To this end, without loss of generality, we fix household j and consider her budget constraint (equation (2.2)) in some period t where the real interest rate is r^{MP}_t and consumption taxes and labor taxes as well as the government debt level are at their steady state levels $\bar{\tau}^C, \bar{\tau}^C, \bar{B}$, respectively. Consistent with our definition of the policy-induced redistribution, we set $(c_{j,t}, l_{j,t}, \frac{b_{j,t}}{1+\tau^C_t}, d_t, w_t) = (\bar{c}_j, \bar{l}_j, \frac{\bar{b}'_j}{1+\bar{\tau}}, \bar{d}, \bar{w})$. This yields the following expression for her policy-induced redistribution:

$$\Xi^{MP}_{j,t} = -\Big(1+\bar{\tau}^C\Big)\bar{c}_j - \bar{b}'_j + (1+r^{MP}_t)\bar{b}_j + \Big(1-\bar{\tau}^L\Big)\bar{w}z_{j,t}\bar{l}_j + \bar{D} + \tilde{Tr}^{MP}_t. \tag{B.9}$$

Solving the government budget constraint with the interest rate at period t (but with constant behavior of the agents) gives $\tilde{Tr}^{MP}_t = -r^{MP}_t\bar{B} + \bar{T}$. This can be interpreted as the policy-induced partial equilibrium transfer. Hence, $\Xi^{MP}_{j,t}$ is only affected by the changed return on savings and the direct effect of the real interest rate on transfers. Inserting $\tilde{Tr}_t$ in equation (B.9) and using the steady state budget constraint of household j yields:

$$\begin{aligned}\Xi^{MP}_{j,t} &= \Big(-r^{MP}_t\Big)\bar{B} + \bar{T} - \Big(-\bar{r}\Big)\bar{B} - \bar{T} + \bar{b}_j\Big(r^{MP}_t - \bar{r}\Big)\\ &= \bar{B}\Big(\bar{r} - r^{MP}_t\Big) - \bar{b}_j\Big(\bar{r} - r^{MP}_t\Big).\end{aligned}$$

Policy-exposure to HANK-UFP. We now derive the policy-exposure with HANK-UFP, $\Xi^{UFP}_{h,t}$. To this end, without loss of generality, we fix household j and consider her budget constraint (equation (2.2)) in some period t, where the real interest rate is at its steady state level $\bar{r}$ and consumption taxes and labor taxes as well as the government debt level are set according to Proposition 1. Consistent with our definition of the policy-induced redistribution, we set $(c_{j,t}, l_{j,t}, \frac{b_{j,t+1}}{1+\tau^C_t}, d_t, w_t) = (\bar{c}_j, \bar{l}_j, \frac{\bar{b}'_j}{1+\bar{\tau}}, \bar{d}, \bar{w})$. This gives the following expression for her policy-induced redistribution:

$$\Xi^{UFP}_{j,t} = -(1+\tau^{C,UFP}_t)\bar{c}_j - \frac{1+\tau^{C,UFP}_t}{1+\bar{\tau}^{C,UFP}}\bar{b}'_j + (1+\bar{r})\frac{1+\tau^C_{t-1}}{1+\bar{\tau}^C}\bar{b}_j + (1-\tau^{L,UFP}_t)\bar{w}z_{j,t}\bar{l}_j + \bar{D} + \tilde{Tr}^{UFP}_t.$$

Dividing by gross consumption taxes, inserting the budget constraint in the original steady state, and using condition (2.21) yields:

$$\frac{\Xi^{UFP}_{j,t}}{1+\tau^{C,UFP}_t} = -(1+\bar{r})\frac{\bar{b}_j}{(1+\bar{\tau}^C)} + (1+\bar{r})\frac{\frac{1+\tau^{C,UFP}_{t-1}}{1+\bar{\tau}^C}\bar{b}_j}{(1+\tau^{C,UFP}_t)} + \frac{\bar{D}}{1+\tau^{C,UFP}_t} - \frac{\bar{D}}{1+\bar{\tau}^C} + \frac{\tilde{Tr}^{UFP}_t}{1+\tau^{C,UFP}_t} - \frac{\bar{Tr}}{1+\bar{\tau}^C}.$$

Rearranging and using condition (2.20) yields:

$$\frac{\Xi^{UFP}_{j,t}}{1+\tau^{C,UFP}_t} = \frac{\bar{b}_j}{(1+\bar{\tau}^C)}(r^{MP}_t - \bar{r}) + D\left(\frac{\bar{1}}{1+\tau^{C,UFP}_t} - \frac{\bar{1}}{1+\bar{\tau}^C}\right) + \frac{\tilde{Tr}^{UFP}_t}{1+\tau^{C,UFP}_t} - \frac{\bar{Tr}}{1+\bar{\tau}^C}.$$

Solving the government budget constraint and using condition (2.22) gives the policy-induced transfer in the HANK-UFP case, $\tilde{Tr}_t^{UFP} = \bar{B}\frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C} - (1+\bar{r})\frac{1+\tau_{t-1}^{C,UFP}}{1+\bar{\tau}^C}\bar{B} + \tilde{T}_t^{UFP}$, where $\tilde{T}_t^{UFP}$ is the policy-induced partial equilibrium tax income, i.e., the tax income with steady state consumption and labor supply but with HANK-UFP tax rates. Inserting this and the steady state transfer as well as rearranging yields:

$$\frac{\Xi_{j,t}^{UFP}}{1+\tau_t^{C,UFP}} = \frac{\bar{b}_j}{(1+\bar{\tau}^C)}(r_t^{MP} - \bar{r}) + D\left(\frac{1}{1+\tau_t^{C,UFP}} - \frac{1}{1+\bar{\tau}^C}\right)$$
$$-\frac{\bar{B}}{(1+\bar{\tau}^C)}(r_t^{MP} - \bar{r}) + \frac{\tilde{T}_t^{UFP}}{1+\tau_t^{C,UFP}} - \frac{\bar{T}}{1+\bar{\tau}^C}.$$

Multiplying with $1+\bar{\tau}^C$ and further rearranging yields:

$$\Xi_{j,t}^{UFP}\frac{1+\bar{\tau}^C}{1+\tau_t^{C,UFP}} = \bar{B}(\bar{r} - r_t^{MP}) - \bar{b}_j(\bar{r} - r_t^{MP})$$
$$+D\left(\frac{1+\bar{\tau}^C}{1+\tau_t^{C,UFP}} - 1\right) + \frac{1+\bar{\tau}^C}{1+\tau_t^{C,UFP}}\tilde{T}_t^{UFP} - \bar{T}.$$

Given condition (2.21), the last three terms add up to zero as shown in Appendix B.1 starting from equation (B.6).[102] Hence,

$$\Xi_{j,t}^{UFP} = \frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C}\Xi_{j,t}^{MP}.$$

B.3 Extension to Sticky Wages

We here follow Auclert et al. (ming). Labor hours are determined by union labor demand and we assume that every worker h provides l_{hkt} hours of work to each continuum of unions indexed $k \in [0,1]$. Total labor effort for person h is therefore $l_{ht} \equiv \int_k l_{hkt} dk$. Each union k aggregates efficient units of work into a union-specific task $L_{ht} = \int z_{ht} l_{hkt} dh$. A competitive labor packer then packages these tasks into aggregate employment services using the technology with constant elasticity of substitution $L_t = \left(\int_k L_{kt}^{\frac{\epsilon-1}{\epsilon}} dk\right)^{\frac{\epsilon}{\epsilon-1}}$ and sells these services to final goods firms at price w_t.

There is a quadratic utility cost of adjusting nominal wage W_{kt} set by union k through an extra additive disutility term $\frac{\nu}{2}\int_k \left(\frac{w_{kt}}{w_{kt-1}} - 1\right)^2 dk$ in household utility (C.33). In each period t, union k sets a common wage w_{kt} per efficient unit for each of its members, and calls upon its members to supply hours according to a uniform rule, such that $l_{hkt} = L_{kt}$. The union sets w_{kt} to maximize the average utility of its members given this allocation rule.

In this setup, all unions choose to set the same wage $w_{kt} = w_t$ at time t and all households work the same number of hours, equal to $l_{ht} = L_t$ so efficiency-weighted hours

[102] Too see this, replace $C_t^*, L_t^*, w_t^*, D_t^*$ in the equations following equation (B.6) with their steady state values.

worked $\int z_{it} l_{it} di$ are also equal to aggregate labor demand L_t. Real wages evolve according to $\frac{1+\pi_t^W}{1+\pi_t} = \frac{w_t}{w_{t-1}}$, where π_t^W is the nominal wage inflation which evolves according to an aggregate non-linear wage Phillips curve

$$\pi_t^w(1+\pi_t^w) = \frac{\epsilon_w}{\nu} \int L_t \left(L_t^{\psi} - \frac{\epsilon_w - 1}{\epsilon_w} \frac{1-\tau_t^L}{1+\tau_t^C} w_t z_{ht} c_{ht}^{-\gamma} \right) dh + \beta \pi_{t+1}^w (1+\pi_{t+1}^w),$$

where ϵ_w denotes the elasticity of substitution among unions.

B.4 Extension to Investment

In this section, we present our extension to investment.

B.4.1 Households

The household problem is the same as the one described in Section (2.2.1) except that now households face the budget constraints described in equation (2.31). In the case in which capital and bonds are perfect substitutes, $\omega_{h,t} = 0$.

B.4.2 Intermediate Good Firms

Intermediate goods are produced by a continuum of intermediate good firms in monopolistically competitive markets. They now produce according to the following Cobb-Douglas production function:

$$y_{j,t} = n_{j,t}^{1-\alpha} k_{j,t}^{\alpha},$$

where $0 < \alpha < 1$, $n_{j,t}$ is labor services, and $k_{j,t}$ is capital services rented in perfectly competitive factor markets. We assume that the intermediate good firms are subject to the same Calvo-pricing as in our baseline model.

B.4.3 Fiscal Policy

The government's budget constraint is given by equation (2.30). In addition, we set $\bar{\tau}^F = 0$. In the case in which bonds and capital are perfect substitutes, $\Omega_t = 0 \; \forall t$.

B.4.4 Equilibrium

Our definition of an equilibrium of this extended economy is analogous to Section 2.2.4.

B.4.5 Proof of equivalence when bonds and capital are perfect substitutes

We now prove that conditions (2.28) and (2.26) together with conditions (2.20) and (2.21) are sufficient conditions for HANK-UFP to generate the same allocation as monetary policy in our HANK model with capital. As in Appendix B.1, the same allocation implies that households' savings in the HANK-UFP case and in the monetary policy case are the same in consumption value terms.

Satisfying the firms' first-order condition. Given that capital subsidies are set according to condition (2.26), the firms' first-order condition (2.25) is the same with HANK-UFP and monetary policy. Note that this also implies that firms' costs are the same since they pay the same net rental rate for the capital they use in the production.

Satisfying household's first-order conditions. Both HANK-UFP and monetary policy generate the same effects on households' first-order conditions since the logic of Appendix B.1 carries over. Hence, if $x_j^{MP} = \{c_{j,t}^*, l_{j,t}^*\}_{t=0}^{\infty}$, satisfies her sequences of Euler equations and labor leisure equations in the monetary policy experiment, $x_j^{UFP} = \{c_{j,t}^*, l_{j,t}^*\}_{t=0}^{\infty}$ does so in the HANK-UFP case.

Satisfying household's budget constraint. The budget constraint of households is now given by equation (2.27). Denote by $a_{h,t} = b_{h,t} + k_{h,t}$ the total savings of a household. As in Appendix B.1, we show that if $x_j^{MP} = \{a_{j,t+1}^*, c_{j,t}^*, l_{j,t}^*\}_{t=0}^{\infty}$ satisfies the sequence of budget constraints of a given household j in the monetary policy experiment, $x_j^{UFP} = \{\frac{1+\tau_t}{1+\bar{\tau}} a_{j,t+1}^*, c_{j,t}^*, l_{j,t}^*\}_{t=0}^{\infty}$ does so with HANK-UFP. First note that the labor income term and the lump-sum income term do not change compared to our baseline result given that the transfers resulting as a residual from the government budget constraint generate the transfer path given by equation (2.23) which we show below to be the case.

The asset income term, however, is now different. It is the same iff:

$$\frac{1+r_t^{MP}}{1+\bar{\tau}^C} a_{j,t}^* - \frac{a_{j,t+1}^*}{1+\bar{\tau}^C} = \frac{1+\bar{r}}{1+\tau_t^{C,UFP}} a_{j,t}^{UFP} - \frac{a_{j,t+1}^{UFP}}{1+\tau_t^{C,UFP}}. \tag{B.10}$$

Hence, we obtain equivalence iff:

$$a_{j,t}^* \frac{1+r_t^{MP}}{1+\bar{\tau}^C} - \frac{a_{j,t+1}^*}{1+\bar{\tau}^C} = \frac{1+\bar{r}}{1+\tau_t^{C,UFP}} \frac{1+\tau_{t-1}^{C,UFP}}{1+\bar{\tau}^C} a_{j,t}^* - \frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C} \frac{a_{j,t+1}^*}{1+\tau_t^{C,UFP}},$$

which holds given Condition (2.20).

Note that $a_{j,t+1}^{UFP} = \frac{1+\tau_t^{C,UFP}}{1+\bar{\tau}^C} a_{j,t+1}^*$ is feasible for every j given condition (2.28) which can be seen by using condition (2.20) in condition (2.28).

Equivalence in government budget constraint. We now show that the residual transfers in the HANK-UFP case indeed follow (2.23). We do so by plugging (2.23) into the government budget constraint in the HANK-UFP case and show that it holds. To make the following calculations more readable, we now denote $K_t^{MP} = K_t^{UFP} = K_t$, $\tau_t^{C,UFP} = \tau_t^C$, $\tau_t^{L,UFP} = \tau_t^L$, $\tau_t^{F,UFP} = \tau_t^F$, $B_t^{UFP} = B_t$, and $r_t^{MP} = r_t$.

The government budget constraint in the HANK-UFP case is given by

$$Tr_t^{UFP} + (1+\bar{r})B_t + \tau_t^F K_t = B_{t+1} + T_t^{UFP}. \tag{B.11}$$

We first plug the transfer path in the HANK-UFP case into (B.11), given by:

$$Tr_t^{UFP} = \frac{1+\tau_t^C}{1+\bar{\tau}^C} Tr_t^{MP} + D_t\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C} - 1\right). \tag{B.12}$$

We then use the government budget constraint in the monetary policy case given by $Tr_t^{MP} = -(1+r_t)\bar{B} + \bar{B} + T_t^{MP} = -r\bar{B} + T_t^{MP}$, $T_t^{MP} = \bar{\tau}^C(Y_t - K_{t+1} + (1-\delta)K_t) + \bar{\tau}^L(Y_t - (r_t+\delta)K_t - D_t)$, $T_t^{UFP} = \tau_t{}^C(Y_t - K_{t+1} + (1-\delta)K_t) + \tau_t^L(Y_t - (r_t+\delta)K_t - D_t)$, $\tau_t^F = \bar{r} - r_t$ (assuming that $\bar{\tau}^F = 0$ and using that $r_t = r_t^k$) and rewrite condition (2.28) by using condition (2.20) into $B_{t+1} = \frac{(1+\tau_{t+1}^C)(1+r_{t+1})}{(1+\bar{\tau}^C)(1+\bar{r})}\bar{B} + \frac{(1+\tau_{t+1}^C)(1+r_{t+1})}{(1+\bar{\tau}^C)(1+\bar{r})}K_{t+1} - K_{t+1}$ which yields

$$\begin{aligned}
\frac{1+\tau_t^C}{1+\bar{\tau}^C}\Big(-r_t\bar{B} + \bar{\tau}^C(Y_t - K_{t+1} + (1-\delta)K_t) + \bar{\tau}^L(Y_t - (r_t+\delta)K_t - D_t)\Big) + D_t\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C} - 1\right) + \\
(\bar{r} - r_t)K_t + \frac{(1+\tau_t^C)(1+r_t)}{(1+\bar{\tau}^C)}\bar{B} + \frac{(1+\tau_t^C)(1+r_t)}{(1+\bar{\tau}^C)}K_t - (1+\bar{r})K_t \\
= \frac{(1+\tau_{t+1}^C)(1+r_{t+1})}{(1+\bar{\tau}^C)(1+\bar{r})}\bar{B} + \frac{(1+\tau_{t+1}^C)(1+r_{t+1})}{(1+\bar{\tau}^C)(1+\bar{r})}K_{t+1} - K_{t+1} \\
+ \tau_t{}^C(Y_t - K_{t+1} + (1-\delta)K_t) + \tau_t^L(Y_t - (r_t+\delta)K_t - D_t)
\end{aligned} \tag{B.13}$$

We will show that this holds by collecting appropriate terms one by one and show that these terms drop out. We start by collecting all terms involving dividends and show that these cancel out:

$$\begin{aligned}
\left(-\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L + \frac{1+\tau_t^C}{1+\bar{\tau}^C} - 1 + \tau_t^L\right)D_t = 0 \\
\iff (-\bar{\tau}^L - \bar{\tau}^L\tau_t^C + 1 + \tau_t^C - 1 - \bar{\tau}^C + \tau_t{}^L + \bar{\tau}^C\tau_t^L)D_t = 0 \\
\iff (\underbrace{-\bar{\tau}^L - \bar{\tau}^C - \bar{\tau}^L\tau_t^C + \tau_t^C + \tau_t{}^L + \bar{\tau}^C\tau_t^L}_{=0,\text{ see (B.7)}})D_t = 0
\end{aligned}$$

Next, we show that all the terms involving $\bar{B}$ are 0.

$$
\begin{aligned}
-r_t \frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{B} + \frac{1+\tau_t^C}{1+\bar{\tau}^C}(1+r_t)\bar{B} - \frac{1+\tau_{t+1}^C}{1+\bar{\tau}^C}\frac{1+r_{t+1}}{1+\bar{r}}\bar{B} = 0 \\
\iff \frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{B} - \frac{1+\tau_t^C}{1+\bar{\tau}^C}\frac{1+\bar{r}}{1+\bar{r}}\bar{B} = 0
\end{aligned}
$$

Hence, equation (B.13) can then be simplified to:

$$
\begin{aligned}
\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C(Y_t - K_{t+1} + (1-\delta)K_t) + (\bar{r} - r_t)K_t + \\
+\frac{(1+\tau_t^C)(1+r_t)}{(1+\bar{\tau}^C)}K_t - (1+\bar{r})K_t \\
= \frac{(1+\tau_{t+1}^C)(1+r_{t+1})}{(1+\bar{\tau}^C)(1+\bar{r})}K_{t+1} - K_{t+1} - \frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L(Y_t - (r_t+\delta)K_t) + \\
{\tau_t}^C(Y_t - K_{t+1} + (1-\delta)K_t) + \tau_t^L(Y_t - (r_t+\delta)K_t)
\end{aligned} \tag{B.14}
$$

Rearranging yields

$$
\begin{aligned}
\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C - \tau_t^C\right)(Y_t - K_{t+1} + (1-\delta)K_t) \\
+\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L - \tau_t^L\right)(Y_t - (r_t+\delta)K_t) = -(\bar{r} - r_t)K_t - \frac{(1+\tau_t^C)(1+r_t)}{(1+\bar{\tau}^C)}K_t + (1+\bar{r})K_t \\
+\frac{1+\tau_t^C}{1+\bar{\tau}^C}K_{t+1} - K_{t+1}
\end{aligned} \tag{B.15}
$$

We next show that all the terms in (B.15) involving Y_t are 0:

$$
\begin{aligned}
\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C Y_t - \tau_t^C Y_t + \frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L Y_t - \tau_t^L Y_t = 0 \\
\iff (1+\tau_t^C)\bar{\tau}^C Y_t - (1+\bar{\tau}^C)\tau_t^C Y_t + (1+\tau_t^C)\bar{\tau}^L Y_t - (1+\bar{\tau}^C)\tau_t^L Y_t = 0 \\
\iff Y_t(\bar{\tau}^C + \bar{\tau}^C\tau_t^C - \tau_t^C - \bar{\tau}^C\tau_t^C + \bar{\tau}^L + \tau_t^C\bar{\tau}^L - \tau_t^L - \bar{\tau}^C\tau_t^L) = 0 \\
\iff Y_t(\underbrace{\bar{\tau}^C - \tau_t^C - \tau_t^L - \bar{\tau}^C\tau_t^L + \bar{\tau}^L + \tau_t^C\bar{\tau}^L}_{=0\text{, see (B.7)}}) = 0
\end{aligned}
$$

We next show that all the terms in (B.15) involving K_{t+1} are 0:

$$
\begin{aligned}
\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C-\tau_t^C\right)K_{t+1}+\frac{1+\tau_t^C}{1+\bar{\tau}^C}K_{t+1}-K_{t+1}&=0\\
\iff \frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C K_{t+1}-\tau_t^C K_{t+1}+\frac{1+\tau_t^C}{1+\bar{\tau}^C}K_{t+1}-K_{t+1}&=0\\
\iff K_{t+1}\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C-(\tau_t^C+1)+\frac{1+\tau_t^C}{1+\bar{\tau}^C}\right)&=0\\
\iff K_{t+1}(\bar{\tau}^C+\tau_t^C\bar{\tau}^C-\tau_t^C-1-\bar{\tau}^C\tau_t^C-\bar{\tau}^C+1+\tau_t^C)&=0
\end{aligned}
$$

We next show that all the terms in (B.15) involving δK_t are 0:

$$
\begin{aligned}
-\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C-\tau_t^C\right)\delta K_t-\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L-\tau_t^L\right)\delta K_t&=0\\
\iff \delta K_t\left(-\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C+\tau_t^C-\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L+\tau_t^L\right)&=0\\
\iff \delta K_t(-(1+\tau_t^C)\bar{\tau}^C+\tau_t^C+\bar{\tau}^C\tau_t^C-(1+\tau_t^C)\bar{\tau}^L+\tau_t^L\bar{\tau}^C+\tau_t^L)&=0\\
\iff \delta K_t(-\bar{\tau}^C-\tau_t^C\bar{\tau}^C+\tau_t^C+\bar{\tau}^C\tau_t^C-\bar{\tau}^L-\tau_t^C\bar{\tau}^L+\tau_t^L\bar{\tau}^C+\tau_t^L)&=0\\
\iff \delta K_t(\underbrace{\tau_t^C+\tau_t^L-\bar{\tau}^C-\bar{\tau}^L-\tau_t^C\bar{\tau}^L+\tau_t^L\bar{\tau}^C}_{=0,\text{ see (B.7)}})&=0
\end{aligned}
$$

Hence, equation (B.15) can now be simplified to:

$$
\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C-\tau_t^C\right)K_t-\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L-\tau_t^L\right)r_tK_t-r_tK_t-K_t+\frac{1+\tau_t^C}{1+\bar{\tau}^C}(1+r_t)K_t=0 \tag{B.16}
$$

We next show that all the terms in (B.16) involving r_tK_t are 0:

$$
\begin{aligned}
-\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L-\tau_t^L\right)r_tK_t-r_tK_t+\frac{1+\tau_t^C}{1+\bar{\tau}^C}r_tK_t&=0\\
\iff r_tK_t\left(-\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L+\tau_t^L-1+\frac{1+\tau_t^C}{1+\bar{\tau}^C}\right)&=0\\
\iff r_tK_t(-\bar{\tau}^L-\tau_t^C\bar{\tau}^L+\tau_t^C+\bar{\tau}^C\tau_t^L-1-\bar{\tau}^C+1+\tau_t^C)&=0\\
\iff r_tK_t(\underbrace{\tau_t^C+\tau_t^L-\bar{\tau}^L-\bar{\tau}^C-\tau_t^C\bar{\tau}^L+\bar{\tau}^C\tau_t^L}_{=0,\text{ see (B.7)}})&=0
\end{aligned}
$$

We finally show that all the terms in (B.16) involving K_t are 0:

$$\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C - \tau_t^C - 1 + \frac{1+\tau_t^C}{1+\bar{\tau}^C}\right)K_t = 0$$

$$K_t(\bar{\tau}^C + \bar{\tau}^C\tau_t^C - \tau_t^C - \bar{\tau}^C\tau_t^C - 1 - \bar{\tau}^C + 1 + \tau_t^C) = 0$$

B.4.6 Proof of equivalence when bonds and capital are imperfect substitutes

When allowing for $r_t \neq r_t^k$, where r_t^k is the return on capital net of depreciation, the budget constraint of a given household j is now given by equation (2.31). We now denote $a_{h,t+1} = k_{h,t+1} + \omega_{h,t+1}$ as the total illiquid asset holdings of a household. To make the following calculations more readable, we now denote $K_t^{MP} = K_t^{UFP} = K_t$, $\tau_t^{C,UFP} = \tau_t^C$, $\tau_t^{L,UFP} = \tau_t^L$, $\tau_t^{F,UFP} = \tau_t^F$, $B_t^{UFP} = B_t$, and $r_t^{MP} = r_t$. Assuming equivalence in illiquid asset income, we obtain an expression for $1 + r_t^{k,UFP}$:

$$\begin{aligned}
\frac{1+r_t^{k,MP}}{1+\bar{\tau}^C} a_{j,t}^* - \frac{a_{j,t+1}^*}{1+\bar{\tau}^C} &= \frac{1+r_t^{k,UFP}}{1+\tau_t^C} \frac{1+\tau_{t-1}^C}{1+\bar{\tau}^C} a_{j,t}^* - \frac{\frac{1+\tau_t^C}{1+\bar{\tau}^C} a_{j,t+1}^*}{1+\tau_t^C} \\
\iff \frac{1+r_t^{k,MP}}{1+\bar{\tau}^C} &= \frac{1+r_t^{k,UFP}}{1+\tau_t^C} \frac{1+\tau_{t-1}^C}{1+\bar{\tau}^C} \\
\iff 1+r_t^{k,MP} &= \left(1+r_t^{k,UFP}\right) \frac{1+\tau_{t-1}^C}{1+\tau_t^C} \\
\iff 1+r_t^{k,UFP} &= \left(1+r_t^{k,MP}\right) \frac{1+\bar{r}}{1+r_t}
\end{aligned}$$

From the firms' first-order conditions, we know:

$$\begin{aligned}
1+r_t^{k,MP} &= 1+r_t^{k,UFP} - \tau_t^F \\
\iff \tau_t^F &= 1+r_t^{k,UFP} - \left(1+r_t^{k,MP}\right) \\
\iff \tau_t^F &= \left(1+r_t^{k,MP}\right) \frac{1+\bar{r}}{1+r_t} - \left(1+r_t^{k,MP}\right) \\
\iff \tau_t^F &= \left(1+r_t^{k,MP}\right) \left(\frac{1+\bar{r}}{1+r_t} - 1\right)
\end{aligned}$$

Equivalence in the government budget constraint. With synthetic capital, Ω_t, as a second illiquid asset, the government budget constraint is given by equation (2.30) and Ω_t is set according to condition (2.32). Doing the same rearrangements as before, this yields:

$$\begin{aligned}
&\frac{1+\tau_t^C}{1+\bar{\tau}^C} \left(-r_t \bar{B} + \bar{\tau}^C (Y_t - K_{t+1} + (1-\delta)K_t) + \bar{\tau}^L (Y_t - (r_t^{k,MP} + \delta)K_t - D_t) \right) \\
&\qquad + D_t \left(\frac{1+\tau_t^C}{1+\bar{\tau}^C} - 1 \right) + (1+r_t^{k,MP}) \left(\frac{1+\bar{r}}{1+r_t} - 1 \right) K_t + \\
&\qquad (1+\bar{r}) \frac{(1+\tau_t^C)(1+r_t)}{(1+\bar{\tau}^C)(1+\bar{r})} \bar{B} + (1+r_t^{k,UFP}) K_t \left(\frac{1+\tau_{t-1}^C}{1+\bar{\tau}^C} - 1 \right) = \\
&\frac{1+\tau_t^C}{1+\bar{\tau}^C} \bar{B} + \left(\frac{1+\tau_t^C}{1+\bar{\tau}^C} - 1 \right) K_{t+1} + \tau_t{}^C (Y_t - K_{t+1} + (1-\delta)K_t) + \tau_t^L (Y_t - (r_t^{k,MP} + \delta)K_t - D_t)
\end{aligned} \tag{B.17}$$

Note that all terms except for the K_t and K_{t+1} terms are the same as in Appendix B.4.5. Hence, all of them cancel out. What remains to be shown, is that the K_t and K_{t+1} terms also cancel out in the model with illiquid capital. We start by showing that the terms in (B.17) involving K_t are 0:

$$\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C(1-\delta)K_t - \frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L\Big(r_t^{k,MP}+\delta\Big)K_t + \Big(1+r_t^{k,MP}\Big)\left(\frac{1+\bar{r}}{1+r_t}-1\right)K_t$$
$$+\Big(1+r_t^{k,UFP}\Big)K_t\left(\frac{(1+\tau_t^C)(1+r_t)}{(1+\bar{\tau}^C)(1+\bar{r})}-1\right) - \tau_t^C(1-\delta)K_t + \tau_t^L\Big(r_t^{k,MP}+\delta\Big)K_t = 0 \quad \text{(B.18)}$$

To do so, we first show that only the terms involving δK_t in (B.18) drop out

$$-\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C\delta K_t - \frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L\delta K_t + \tau_t^C\delta K_t + \tau_t^L\delta K_t = 0$$
$$\delta K_t(-\bar{\tau}^C - \bar{\tau}^C\tau_t^C - \bar{\tau}^L + \tau_t^C\bar{\tau}^L + \tau_t^C + \bar{\tau}^C\tau_t^C + \tau_t^L + \bar{\tau}^C\tau_t^L) = 0$$
$$\delta K_t(\underbrace{-\bar{\tau}^C - \bar{\tau}^L + \tau_t^C\bar{\tau}^L + \tau_t^C + \tau_t^L + \bar{\tau}^C\tau_t^L}_{=0,\text{ see (B.7)}}) = 0$$

We next show that the rest of equation (B.18) is also zero:

$$\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C K_t - \frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L r_t^{k,MP}K_t + \Big(1+r_t^{k,MP}\Big)\frac{1+\bar{r}}{1+r_t}K_t - (1+r_t^{k,MP})K_t$$
$$+\Big(1+r_t^{k,MP}\Big)\frac{1+\bar{r}}{1+r_t}\frac{(1+\tau_t^C)(1+r_t)}{(1+\bar{\tau}^C)(1+\bar{r})}K_t - \Big(1+r_t^{k,MP}\Big)\frac{1+\bar{r}}{1+r_t}K_t$$
$$-\tau_t^C K_t + \tau_t^L r_t^{k,MP}K_t = 0$$

Rearranging yields

$$\Big(1+r_t^{k,MP}\Big)K_t\left(\frac{1+\bar{r}}{1+r_t} - 1 + \frac{1+\tau_t^C}{1+\bar{\tau}^C} - \frac{1+\bar{r}}{1+r_t}\right)$$
$$+\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C K_t - \frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L r_t^{k,MP}K_t - \tau_t^C K_t + \tau_t^L r_t^{k,MP}K_t = 0$$

We now show first that the following terms are 0:

$$\left(\frac{1+\tau_t^C}{1+\bar{\tau}^C}-1\right)K_t + \frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C K_t - \tau_t^C K_t = 0$$
$$\iff (1+\tau_t^C - 1 - \bar{\tau}^C + \bar{\tau}^C + \tau_t^C\bar{\tau}^C - \tau_t^C - \bar{\tau}^C\tau_t^C)K_t = 0$$

We next show that the following terms are 0:

$$
\begin{aligned}
r_t^{k,MP} K_t \left(\frac{1+\tau_t^C}{1+\bar{\tau}^C} - 1 - \frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^L + \tau_t^L \right) &= 0 \\
r_t^{k,MP} K_t (1 + \tau_t^C - 1 - \bar{\tau}^C - \bar{\tau}^L - \tau_t^C \bar{\tau}^L + \tau_t^L + \bar{\tau}^C \tau_t^L) &= 0 \\
r_t^{k,MP} K_t (\underbrace{\tau_t^C + \tau_t^L - \bar{\tau}^C - \bar{\tau}^L - \tau_t^C \bar{\tau}^L + \bar{\tau}^C \tau_t^L}_{=0,\text{ see (B.7)}}) &= 0
\end{aligned}
$$

Finally, we show that all the terms involving K_{t+1} in (B.17) are 0:

$$
\begin{aligned}
-\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C K_{t+1} - \frac{1+\tau_t^C}{1+\bar{\tau}^C} K_{t+1} + K_{t+1} + \tau_t^C K_{t+1} &= 0 \\
K_{t+1} \left(-\frac{1+\tau_t^C}{1+\bar{\tau}^C}\bar{\tau}^C - \frac{1+\tau_t^C}{1+\bar{\tau}^C} + 1 + \tau_t^C \right) &= 0 \\
K_{t+1} (-\bar{\tau}^C - \tau_t^C \bar{\tau}^C - 1 - \tau_t^C + 1 + \bar{\tau}^C + \tau_t^C + \bar{\tau}^C \tau_t^C) &= 0
\end{aligned}
$$

Appendix C

Appendix for Chapter 3

C.1 Equilibrium

We define a linearized perfect-foresight transition economy of the model. That is, we always refer to the linearized versions of the relevant model equations.

Given exogenous paths for the supply and the demand shocks, $\{\epsilon_t, \epsilon_t^*, \xi_t, \xi_t^*\}_{t=0}^{\infty}$, a linearized perfect-foresight equilibrium is a set of aggregates $\{\pi_t, \pi_t^*, \pi_{H,t}, \pi_{H,t}^*, \pi_{F,t}, \pi_{F,t}^*, i_t, i_t^*, r_t, r_t^*, \tau_t, \tau_t^*, w_t, w_t^*, c_t, c_t^*, Y_t, Y_t^*, N_t, N_t^*, \Delta\mathcal{E}_t, s_t, Q_t, A_t, A_t^*, B_t, B_t^*, B_{F,t}\}_{t=0}^{\infty}$ such that:

1. The paths of aggregate consumption at Home and at Foreign $\{c_t, c_t^*\}_{t=0}^{\infty}$ are consistent with the linearized aggregate consumption functions (4), and the path of household asset holdings in Home and in Foreign $\{A_t, A_t^*\}_{t=0}^{\infty}$ are consistent with the budget constraints (2), aggregated across households in Home and in Foreign.

2. The real wages $\{w_t, w_t^*\}_{t=0}^{\infty}$ are consistent with (8) and the counterpart in Foreign.

3. The paths of $\{N_t, N_t^*, Y_t, Y_t^*\}_{t=0}^{\infty}$ satisfy the aggregate production functions in Home and in Foreign (7).

4. The paths of $\{\pi_{H,t}, \pi_{F,t}^*, Y_t, Y_t^*, s_t\}_{t=0}^{\infty}$ are consistent with the national Philips curves (10).

5. The paths of $\{\pi_{H,t}^*, \pi_{F,t}\}_{t=0}^{\infty}$ are consistent with the law of one price stated in the main text.

6. Nominal interest rates and the change in the nominal exchange rate satisfy the interest rate rules given above and the UIP (11) condition holds.

7. The evolution of the government debt levels and taxes $\{B_t, B_t^*, \tau_t, \tau_t^*\}_{t=0}^{\infty}$ are consistent with the government budget constraints (16) and the feedback function for taxes (17).

8. CPI rates in both countries $\{\pi_t, \pi_t^*\}_{t=0}^{\infty}$ are consistent with the definition of the CPI given by (6).

9. The net foreign asset position $\{B_{F,t}\}_{t=0}^{\infty}$ evolves according to the home budget constraint (20).

10. Terms of trade $\{s_t\}_{t=0}^{\infty}$ and the real exchange rate $\{Q_t\}_{t=0}^{\infty}$ evolve as defined in the main text.

11. The bond markets (19), and the goods markets clear (21).

C.2 Derivations and Proofs

C.2.1 Deriving the Canonical Form

We now derive the canonical form of the model, that is the model in terms of union-wide variables. To this end, we define union-wide variables, as e.g., union-wide GDP, as $\widehat{X_t^W} = n\widehat{X_t} + (1-n)\widehat{X_t^*}$.

Using the goods market clearing conditions, we have,

$$\widehat{y_t^W} = n\widehat{C_{H,t}} + (1-n)\widehat{C_{F,t}} = \widehat{C_t^W}. \tag{C.1}$$

From national aggregate consumption functions to union-wide IS equation. The aggregate consumption functions of the countries are:

$$\begin{aligned} \boldsymbol{c} &= \mathcal{C}(\boldsymbol{w}, \boldsymbol{N}, \boldsymbol{i}, \boldsymbol{\pi}, \boldsymbol{\tau}) \\ \boldsymbol{c}^* &= \mathcal{C}^*(\boldsymbol{w}^*, \boldsymbol{N}^*, \boldsymbol{i}^*, \boldsymbol{\pi}^*, \boldsymbol{\tau}^*) \end{aligned} \tag{C.2}$$

Linearizing these consumption functions around the deterministic steady state:

$$\begin{aligned} \widehat{\boldsymbol{c}} &= \mathcal{C}_w\widehat{\boldsymbol{w}} + \mathcal{C}_N\widehat{\boldsymbol{N}} + \mathcal{C}_i\widehat{\boldsymbol{i}} + \mathcal{C}_\pi\widehat{\boldsymbol{\pi}} + \mathcal{C}_\tau\widehat{\boldsymbol{\tau}} \\ \widehat{\boldsymbol{c}^*} &= \mathcal{C}^*_{w^*}\widehat{\boldsymbol{w}^*} + \mathcal{C}^*_{N^*}\widehat{\boldsymbol{N}^*} + {\mathcal{C}^*}_{i^*}\widehat{\boldsymbol{i}^*} + {\mathcal{C}^*}_{\pi^*}\widehat{\boldsymbol{\pi}^*} + {\mathcal{C}^*}_{\tau^*}\widehat{\boldsymbol{\tau}^*}. \end{aligned} \tag{C.3}$$

Given symmetric countries, we have $\mathcal{C}(\boldsymbol{w}, \boldsymbol{N}, \boldsymbol{i}, \boldsymbol{\pi}, \boldsymbol{\tau}) = \mathcal{C}^*(\boldsymbol{w}^*, \boldsymbol{N}^*, \boldsymbol{i}^*, \boldsymbol{\pi}^*, \boldsymbol{\tau}^*)$ and thus, $(\mathcal{C}_w, \mathcal{C}_N, \mathcal{C}_i, \mathcal{C}_\pi, \mathcal{C}_\tau) = ({\mathcal{C}^*}_{w^*}, {\mathcal{C}^*}_{N^*}, {\mathcal{C}^*}_{i^*}, {\mathcal{C}^*}_{\pi^*}, {\mathcal{C}^*}_{\tau^*})$. Using this, we can write world consumption as:

$$\begin{aligned} \widehat{\boldsymbol{c}^{\boldsymbol{W}}} &= n\widehat{\boldsymbol{c}} + (1-n)\widehat{\boldsymbol{c}^*} \\ &= \mathcal{C}_w(n\widehat{\boldsymbol{w}} + (1-n)\widehat{\boldsymbol{w}^*}) + \mathcal{C}_N(n\widehat{\boldsymbol{N}} + (1-n)\widehat{\boldsymbol{N}^*}) \\ &\quad + \mathcal{C}_i(n\widehat{\boldsymbol{i}} + (1-n)\widehat{\boldsymbol{i}^*}) + \mathcal{C}_\pi(n\widehat{\boldsymbol{\pi}} + (1-n)\widehat{\boldsymbol{\pi}^*}) + \mathcal{C}_\tau(n\widehat{\boldsymbol{\tau}} + (1-n)\widehat{\boldsymbol{\tau}^*}) \\ &= \mathcal{C}_w\widehat{\boldsymbol{w}^{\boldsymbol{W}}} + \mathcal{C}_N\widehat{\boldsymbol{N}^{\boldsymbol{W}}} + \mathcal{C}_i\widehat{\boldsymbol{i}^{\boldsymbol{W}}} + \mathcal{C}_\pi\widehat{\boldsymbol{\pi}^{\boldsymbol{W}}} + \mathcal{C}_\tau\widehat{\boldsymbol{\tau}^{\boldsymbol{W}}} \end{aligned} \tag{C.4}$$

Using the linearized version of (8)

$$\hat{w}_t^H = -\alpha_H(1-n)\hat{s}_t^H$$

and its Foreign country counterpart

$$\hat{w}_t^F = -\alpha_H n \hat{s}_t^F$$

it follows that $\widehat{\boldsymbol{w}^{\boldsymbol{W}}} = n\widehat{\boldsymbol{w}^{\boldsymbol{H}}} + (1-n)\widehat{\boldsymbol{w}^{\boldsymbol{F}}} = 0$ because for terms of trade, s, $\widehat{\boldsymbol{s}}^H = -\widehat{\boldsymbol{s}}^F$ holds. Using this and $\widehat{\boldsymbol{c}^{\boldsymbol{W}}} = \widehat{\boldsymbol{y}^{\boldsymbol{W}}} = \widehat{\boldsymbol{N}^{\boldsymbol{W}}}$ (see (C.1) and (7)) and writing $\mathcal{C}_y = \mathcal{C}_N$, we obtain:

$$\widehat{\boldsymbol{y}^{\boldsymbol{W}}} = \mathcal{C}_y \widehat{\boldsymbol{y}^{\boldsymbol{W}}} + \mathcal{C}_i \widehat{\boldsymbol{i}^{\boldsymbol{W}}} + \mathcal{C}_\pi \widehat{\boldsymbol{\pi}^{\boldsymbol{W}}} + \mathcal{C}_\tau \widehat{\boldsymbol{\tau}^{\boldsymbol{W}}} \tag{C.5}$$

Furthermore, we can use national government budget constraints and tax feedback functions to solve for taxes:

$$\begin{aligned} \widehat{\tau}_{t+1} &= ((1+\bar{i}) - \gamma\frac{\bar{\tau}}{b_y})\widehat{\tau}_t + \bar{i}\widehat{i}_t - (1+\bar{i})\widehat{\pi}_{t+1} - \gamma\frac{\bar{\tau}}{b_y}\widehat{y}_t \\ \widehat{\tau^*}_{t+1} &= ((1+\bar{i}) - \gamma\frac{\bar{\tau}}{b_y})\widehat{\tau}_t + \bar{i}\widehat{i^*}_t - (1+\bar{i})\widehat{\pi^*}_{t+1} - \gamma\frac{\bar{\tau}}{b_y}\widehat{y^*}_t, \end{aligned} \tag{C.6}$$

and aggregating gives:

$$\widehat{\tau^W}_{t+1} = ((1+\bar{i}) - \gamma\frac{\bar{\tau}}{b_y})\widehat{\tau}_t + \bar{i}\widehat{i^W}_t - (1+\bar{i})\widehat{\pi^W}_{t+1} - \gamma\frac{\bar{\tau}}{b_y}\widehat{y^W}_t \tag{C.7}$$

Hence, we can stack taxes as:

$$\begin{aligned} \widehat{\boldsymbol{\tau}} = \tau(\widehat{\boldsymbol{y}}, \widehat{\boldsymbol{i}}, \widehat{\boldsymbol{\pi}}),\ \widehat{\boldsymbol{\tau}^*} &= \tau(\widehat{\boldsymbol{y}^*}, \widehat{\boldsymbol{i}^*}, \widehat{\boldsymbol{\pi}^*}) \\ \widehat{\boldsymbol{\tau}^{\boldsymbol{W}}} &= \tau(\widehat{\boldsymbol{y}^{\boldsymbol{W}}}, \widehat{\boldsymbol{i}^{\boldsymbol{W}}}, \widehat{\boldsymbol{\pi}^{\boldsymbol{W}}}) \end{aligned} \tag{C.8}$$

Using this, we can write:

$$\widehat{\boldsymbol{y}^{\boldsymbol{W}}} = \mathcal{C}_y \widehat{\boldsymbol{y}^{\boldsymbol{W}}} + \mathcal{C}_i \widehat{\boldsymbol{i}^{\boldsymbol{W}}} + \mathcal{C}_\pi \widehat{\boldsymbol{\pi}^{\boldsymbol{W}}} + \mathcal{C}_\tau \widehat{\boldsymbol{\tau}^{\boldsymbol{W}}}(\widehat{\boldsymbol{y}^{\boldsymbol{W}}}, \widehat{\boldsymbol{i}^{\boldsymbol{W}}}, \widehat{\boldsymbol{\pi}^{\boldsymbol{W}}}) \tag{C.9}$$

and:

$$\widehat{\boldsymbol{y}^{\boldsymbol{W}}} = \underbrace{[\mathcal{C}_y + \mathcal{C}_\tau \mathcal{T}_y]}_{\tilde{\mathcal{C}}_y} \widehat{\boldsymbol{y}^{\boldsymbol{W}}} + \underbrace{[\mathcal{C}_i + \mathcal{C}_\tau \mathcal{T}_i]}_{\tilde{\mathcal{C}}_i} \widehat{\boldsymbol{i}^{\boldsymbol{W}}} + \underbrace{[\mathcal{C}_\pi + \mathcal{C}_\tau \mathcal{T}_\pi]}_{\tilde{\mathcal{C}}_\pi} \widehat{\boldsymbol{\pi}^{\boldsymbol{W}}} \tag{C.10}$$

where $\mathcal{T}_y$, $\mathcal{T}_r$ and $\mathcal{T}_\pi$ are derivative matrices for the maps $\widehat{\boldsymbol{\tau}}(\widehat{\boldsymbol{y}}, \widehat{\boldsymbol{i}}, \widehat{\boldsymbol{\pi}})$.

Importantly, the world IS equation (C.10) is potentially only affected by different exchange rate regimes through its effects on the world interest rate.

Philips Curve. Using the consumer price indexes and the law of one prices, one obtains:

$$\pi_t^W = n\pi_{H,t} + (1-n)\pi_{F,t} \tag{C.11}$$

Hence, world CPI inflation is just the weighted average of domestic producer price inflation. Using this and aggregating the two national Philips curves gives the union-wide Philips curve:

$$\pi_t^W = \beta\pi_{t+1}^W + \kappa^W \hat{y_t}^W, \tag{C.12}$$

with $\kappa^W = \kappa(\phi + \gamma)$.

Stacking the World Phillips Curve yields:

$$\Pi_\pi \widehat{\boldsymbol{\pi}^{\boldsymbol{W}}} = \Pi_y \widehat{\boldsymbol{y}^{\boldsymbol{W}}}, \tag{C.13}$$

where

$$\Pi_\pi = \begin{pmatrix} 1 & -\beta & 0 & \dots \\ 0 & 1 & -\beta & \dots \\ 0 & 0 & 1 & \dots \\ \vdots & \vdots & \vdots & \ddots \end{pmatrix}, \Pi_y = \kappa^W \omega I, \Pi_y = \kappa^W \sigma^{-1} I \tag{C.14}$$

Lemma 4 *In a two-country model of symmetric countries, we can describe the canonical system for world variables by the following two linear mappings:*

$$\Pi_\pi \widehat{\boldsymbol{\pi}^{\boldsymbol{W}}} = \Pi_y \widehat{\boldsymbol{y}^{\boldsymbol{W}}}, \tag{C.15}$$

$$\widehat{\boldsymbol{y}^{\boldsymbol{W}}} = \tilde{\mathcal{C}}_y \widehat{\boldsymbol{y}^{\boldsymbol{W}}} + \tilde{\mathcal{C}}_i \widehat{\boldsymbol{i}^{\boldsymbol{W}}} + \tilde{\mathcal{C}}_\pi \widehat{\boldsymbol{\pi}^{\boldsymbol{W}}} \tag{C.16}$$

and a function for the nominal interest rate, i_t^W*.*

Note in particular that different exchange rate regimes can only potentially influence world variables through their effect on the world interest rates.

C.3 The medium-sized HANK2

Our medium-sized HANK^2 model extends our baseline model in Section 2 such that its structure in each country mimics the closed-economy set up of Bayer et al. (ming), except for the fact that there is trade across the two countries, both in goods and financial markets. A brief overview of the extensions that we made can be found in Section 4. Here, we present the entire model from scratch.

Each country consists of a firm sector and a household sector. The firm sector of each country comprises (a) perfectly competitive intermediate goods producers who rent out labor services and capital on national labor and a national capital market, respectively; (b) final goods producers that face monopolistic competition when selling differentiated final goods, in turn, produced on the basis of homogeneous domestic intermediate inputs; (c) a representative consumption good bundler bundling domestic and imported foreign final goods to consumption goods; (d) producers of capital goods that turn consumption goods into capital subject to adjustment costs; (e) labor packers that produce labor services combining differentiated labor from (f) unions that differentiate raw labor rented out from households. Price setting for the final goods, as well as wage setting by unions, is subject to a pricing friction à la Calvo (1983). Only final goods are traded across countries.

In each country, there is a continuum of households of size $n \in (0, 1)$ and $1 - n$, respectively, such that the total population is 1. Households in both countries consume a bundle that consists of domestically produced and imported goods. Households earn income from supplying (raw) labor and capital to the national labor and the national capital markets and from owning firms in their respective country. Households absorb all rents that stem from the market power of unions and final good producers, and decreasing returns to scale in capital goods production.

In the baseline, there is a common monetary authority and the exchange rate is permanently fixed. Fiscal policy is run at the country level. It levies taxes on labor income and profits, issues bonds, and adjusts taxes to stabilize the level of outstanding debt in the long run. Public debt is risk-free and thus yields the same return in both countries, in turn, determined by monetary policy by means of a simple interest rate feedback rule. We assume that countries are perfectly symmetric and differ only because of asymmetric shocks and different parameterizations. In what follows, our exposition thus focuses on the domestic economy and uses an asterisk to denote foreign variables whenever they are relevant.

C.3.1 Households

The household sector is subdivided into two types of agents: workers and entrepreneurs. The transition between both types is stochastic. Both rent out physical capital, but only workers supply labor. The efficiency of a worker's labor evolves randomly exposing worker-households to labor-income risk. Entrepreneurs do not work but earn all pure rents in the economy except for the rents of unions which are equally distributed across workers. All households self-insure against the income risks they face by saving in a liquid nominal asset (bonds) and a less liquid asset (capital). Trading illiquid assets is subject to random participation in the capital market. To be specific, there is a continuum of ex-ante identical households of measure n, indexed by i. Households are infinitely lived, have time-separable preferences with time discount factor β, and derive felicity from consumption c_{it} and leisure. They obtain income from supplying labor, n_{it}, from renting out capital, k_{it}, and from earning interest on bonds, b_{it}, and potentially from profits or union transfers. Households pay taxes on labor and profit income.

Productivity, labor supply, and labor income

A household's gross labor income $w_t n_{it} h_{it}$ is composed of the aggregate wage rate on raw labor, w_t, the household's hours worked, n_{it}, and its idiosyncratic labor productivity, h_{it}. We assume that productivity evolves according to a log-AR(1) process with time-varying volatility and a fixed probability of transition between the worker and the entrepreneur state:

$$\tilde{h}_{it} = \begin{cases} \exp(\rho_h \log \tilde{h}_{it-1} + \epsilon^h_{it}) & \text{with probability } 1-\zeta \text{ if } h_{it-1} \neq 0, \\ 1 & \text{with probability } \iota \text{ if } h_{it-1} = 0, \\ 0 & \text{else.} \end{cases} \tag{C.17}$$

with individual productivity $h_{it} = \frac{\tilde{h}_{it}}{\int \tilde{h}_{it} di}$ such that $\tilde{h}_{it}$ is scaled by its cross-sectional average, $\tilde{h}_{it}di$, to make sure that average worker productivity is constant. The shocks ϵ^h_{it} to productivity are normally distributed with variance $\sigma^2_{h,t}$. With probability ζ households become entrepreneurs ($h = 0$). With probability ι an entrepreneur returns to the labor force with median productivity. In our baseline specification, an entrepreneur obtains a share of the pure rents (aside from union rents), Π^F_t, in the economy (from monopolistic competition in the goods sector and the creation of capital). We assume that the claim to the pure rent cannot be traded as an asset. Union rents, Π^U_t are distributed lump sum across workers, leading to labor-income compression. For tractability, we assume union profits to be taxed at a fixed rate independent of the recipient's labor income.[103]

With respect to leisure and consumption, households have Greenwood et al. (1988) (GHH) preferences and maximize the discounted sum of felicity:

$$E_0 \max_{\{c_{it}, n_{it}\}} \sum_{t=0}^{\infty} \beta^t u[c_{it} - G(h_{it}, n_{it})] \tag{C.18}$$

The maximization is subject to the budget constraints described further below. The felicity function u exhibits a constant relative risk aversion (CRRA) with risk aversion parameter $\xi > 0$,

$$u(x_{it}) = \frac{1}{1-\xi} x_{it}^{1-\xi}, \tag{C.19}$$

where $x_{it} = c_{it} - G(h_{it}, n_{it})$ is household i's composite demand for goods consumption c_{it} and leisure and G measures the dis-utility from work. The consumption good c is a bundle of domestic and imported foreign final goods as described in Section C.3.2.

[103]This modeling strategy serves two purposes. First and foremost, it generally solves the problem of the allocation of pure rents without distorting factor returns and without introducing another tradable asset. Second, we use the entrepreneur state in particular – a transitory state in which incomes are very high – to match the income and wealth distribution following the idea by Castaneda et al. (1998). The entrepreneur state does not change the asset returns or investment opportunities available to households.

The household's labor income gets taxed at rate τ_t, such that its net labor income is given by

$$(1-\tau_t)w_t h_{it} n_{it}, \tag{C.20}$$

where w_t is the aggregate wage rate. Given net labor income, the first-order condition for labor supply is

$$\frac{\partial G(h_{it}, n_{it})}{\partial n_{it}} = (1-\tau_t)w_t h_{it} = \frac{y_{it}}{n_{it}}. \tag{C.21}$$

Assuming that G has a constant elasticity w.r.t. n, $\frac{\partial G(h_{it},n_{it})}{n_{it}} = (1+\gamma)\frac{G(h_{it},n_{it})}{n_{it}}$ with $\gamma > 0$, we can simplify the expression for the composite consumption good, x_{it}, making use of this first-order condition (C.21), and substitute $G(h_{it}, n_{it})$ out of the individual planning problem:

$$x_{it} = c_{it} - G(h_{it}, n_{it}) = c_{it} - \frac{1}{1+\gamma} y_{it}. \tag{C.22}$$

When the Frisch elasticity of labor supply is constant and the tax schedule has the form (C.20), the dis-utility of labor is always a fraction of labor income and constant across households. Therefore, in both the household's budget constraint and felicity function, only after-tax income enters and neither hours worked nor productivity appears separately.

What remains to be determined is individual and aggregate effective labor supply. Without further loss of generality, we assume $G(h_{it}, n_{it}) = h_{it}\frac{n_{it}^{1+\gamma}}{1+\gamma}$. This functional form simplifies the household problem in the stationary equilibrium as h_{it} drops out from the first-order condition and all households supply the same number of hours $n_{it} = N(w_t)$. Total effective labor input, $\int n_{it} h_{it} di$, is hence also equal to $N(w_t)$ because we normalized $\int h_{it} di = 1$.[104]

Households also receive profit income from union profits Π_t^U or firms profits Π_t^{fi} as workers or entrepreneurs, respectively. Both profits get taxed at rate τ_t. What is more, households may receive *non-distortionary* targeted transfer as minimum income benefits tr_{it}. All together, after-tax non-capital income, plugging in the optimal supply of hours, is then:

$$y_{it} = [(1-\tau_t)w_t]^{\frac{1+\gamma}{\gamma}} h_{it} + \mathbb{I}_{h_{it}\neq 0}(1-\tau_t)\Pi_t^U + \mathbb{I}_{h_{it}=0}(1-\tau_t)\Pi_t^{fi} + tr_{it}. \tag{C.23}$$

[104]This means that we can read off average productivity risk from the estimated income risk series in the literature. Without scaling the labor dis-utility by productivity, we would need to translate productivity risk to income risk through the endogenous hour response.

Consumption, savings, and portfolio choice

Given this labor income, households optimize inter-temporally subject to their budget constraint:

$$c_{it} + b_{it+1} + q_t k_{it+1} = y_{it} + b_{it}\frac{R(b_{it}, R^b_t)}{\pi^{CPI}_t} + (q_t + r_t)k_{it}k_{it+1} \geq 0,\; b_{it+1} \geq \underline{B} \tag{C.24}$$

b_{it} is real bond holdings, k_{it} is the amount of illiquid assets, q_t is the price of these assets, r_t is their dividend, $\pi^{CPI}_t = \frac{P_t}{P_{t-1}}$ is realized domestic CPI inflation, and R is the gross nominal interest rate on bonds, which depends on the portfolio position of the household and the central bank's interest rate R^b_t, which is set one period before.

All households that do not participate in the capital market ($k_{it+1} = k_{it}$) still obtain dividends and can adjust their bond holdings. Depreciated capital has to be replaced for maintenance, such that the dividend, r_t, is the net return on capital. Holdings of bonds have to be above an exogenous debt limit $\underline{B}$; and holdings of capital have to be non-negative.

Substituting the expression $c_{it} = x_{it} + \frac{1}{1+\gamma}\left[(1-\tau_t)w_t\right]^{\frac{1+\gamma}{\gamma}} h_{it}$ for consumption, we obtain the budget constraint for the composite leisure-consumption good:

$$x_{it} + b_{it+1} + q_t k_{it+1} = b_{it}\frac{R(b_{it}, R^b_t)}{\pi_t} + (q_t + r_t)k_{it} + z_{it}, \quad k_{it+1} \geq 0,\; b_{it+1} \geq \underline{B}, \tag{C.25}$$

where $z_{it} = \frac{\gamma}{1+\gamma}\left[(1-\tau_t)w_t\right]^{\frac{1+\gamma}{\gamma}} h_{it} + \mathbb{I}_{h_{it}\neq 0}(1-\tau_t)\Pi^U_t + \mathbb{I}_{h_{it}=0}(1-\tau_t)\Pi^{fi}_t + tr_{it}$ is income corrected for the dis-utility of labor.

Households make their savings choices and their portfolio choice between liquid bonds and illiquid capital in light of a capital market friction that renders capital illiquid because participation in the capital market is random and i.i.d. in the sense that only a fraction, λ, of households are selected to be able to adjust their capital holdings in a given period. This means that we specify:

$$R(b_{it}, R^b_t) = \begin{cases} R^b_t & \text{if } b_{it} \geq 0 \\ R^b_t + \bar{R} & \text{if } b_{it} < 0 \end{cases}. \tag{C.26}$$

The extra wedge for unsecured borrowing, $\bar{R}$, creates a mass of households with zero unsecured credit but with the possibility to borrow, though at a penalty rate.

Since a household's saving decision——(b'_a, k') for the case of adjustment and (b'_n, k') for non-adjustment—will be some non-linear function of that household's wealth and productivity, inflation and all other prices will be functions of the domestic joint distribution, Θ_t, of (b, k, h) in t and the foreign joint distribution, Θ^*_t. This makes Θ and Θ^* state variables of the household's planning problem and these distributions evolve as a result of the economy's reaction to aggregate shocks. For simplicity, we summarize all effects of aggregate state variables, including the distributions of wealth and income, by writing the dynamic planning problem with time-dependent continuation values.

This leaves us with three functions that characterize the household's problem: value function V^a for the case where the household adjusts its capital holdings, the function V^n for the case in which it does not adjust, and the expected continuation value, $\mathbb{W}$, over both:

$$
\begin{aligned}
V_t^a(b,k,h) &= \max_{k',b'_a} u[x(b,b'_a,k,k',h] + \beta \mathbb{E}_t \mathbb{W}_{t+1}(b'_a,k',h) \\
V_t^n(b,k,h) &= \max_{b'_n} u[x(b,b'_n,k,k,h] + \beta \mathbb{E}_t \mathbb{W}_{t+1}(b'_n,k,h) \\
\mathbb{W}_{t+1}(b',k',h) &= \lambda V_{t+1}^a(b',k',h) + (1-\lambda) V_{t+1}^n(b',k,h).
\end{aligned}
\tag{C.27}
$$

Expectations about the continuation value are taken with respect to all stochastic processes conditional on the current states. Maximization is subject to the corresponding budget constraint.

C.3.2 Firm sector

The firm sector of each country consists of five sub-sectors: (a) a labor sector composed of unions that differentiate raw labor and labor packers who buy differentiated labor and then sell labor services to intermediate goods producers, (b) intermediate goods producers who hire labor services and rent out capital to produce goods, (c) final goods producers who differentiate intermediate goods and then sell them to (d) goods bundlers who bundle them with foreign final goods and finally sell them as consumption goods to households and to (e) capital goods producers, who turn bundled goods into capital goods. None of these products and goods can be traded between both countries, except for the differentiated final goods.

When profit maximization decisions in the firm sector require inter-temporal decisions (i.e. in price and wage setting and in producing capital goods), we assume for tractability that they are delegated to a mass-zero group of households (managers) that are risk-neutral and compensated by a share in profits. They do not participate in any asset market and have the same discount factor as all other households. Since managers are a mass-zero group in the economy, their consumption does not show up in any resource constraint, and all but the unions' profits go to the entrepreneur households (whose $h = 0$). Union profits go lump-sum to worker households.

Labor packers and unions

Worker households sell their labor services to a mass-n_A continuum of unions indexed by j, each of whom offers a different variety of labor to labor packers who then provide labor services to intermediate goods producers. Labor packers produce final labor services according to the production function

$$N_t = \left(\int_0^{n_A} \hat{n}_{jt}^{\frac{\eta_W - 1}{\eta_W}} dj \right)^{\frac{\eta_W}{\eta_W - 1}}. \tag{C.28}$$

out of labor varieties $\hat{n}_{jt}$. Cost minimization by labor packers implies that each variety of labor, each union j, faces a downward-sloping demand curve

$$\hat{n}_{jt} = \left(\frac{W_{jt}}{W_t^{fi}} \right)^{-\eta_W} N_t \tag{C.29}$$

where W_{jt} is the nominal wage set by union j and W_t^{fi} is the nominal wage at which labor packers sell labor services to final goods producers. Since unions have market power, they pay the households a wage lower than the price at which they sell labor to labor packers. Given the nominal wage W_t at which they buy labor from households and given the nominal wage index W_t^{fi}, unions seek to maximize their discounted stream of profits. However, they face a Calvo (1983) type adjustment friction with indexation with the probability λ_w to keep wages constant. They therefore maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \lambda_w^t \frac{W_t^{fi}}{P_t} N_t \left\{ \left(\frac{W_{jt}(\bar{\pi}_W)^t}{W_t^{fi}} - \frac{W_t}{W_t^{fi}} \right) \left(\frac{W_{jt}(\bar{\pi}_W)^t}{W_t^{fi}} \right)^{-\eta_W} \right\}. \tag{C.30}$$

by setting W_{jt} in period t and keeping it constant except for indexation to π_W, the steady state wage inflation rate.

Since all unions are symmetric, we focus on a symmetric equilibrium and obtain the linearized wage Phillips curve from the corresponding first-order condition as follows, leaving out all terms irrelevant at a first-order approximation around the stationary equilibrium:

$$\log \left(\frac{\pi_t^W}{\bar{\pi}^W} \right) = \beta \mathbb{E}_t \log \left(\frac{\pi_{t+1}^W}{\bar{\pi}^W} \right) + \kappa_w \left(mc_t^w - \frac{1}{\mu^W} \right), \tag{C.31}$$

with $\pi_t^W := \frac{W_t^{fi}}{W_{t-1}^{fi}} = \frac{w_t^{fi}}{w_{t-1}^{fi}} \pi_t^{CPI}$ being domestic wage inflation, w_t and w_t^{fi} being the respective *real* wages for households and firms, $mc_t^w = \frac{w_t}{w_t^{fi}}$ is the mark-down of wages the unions pay to households, W_t, relative to the wages charged to firms, W_t^{fi} and $\kappa_w = \frac{(1-\lambda_w)(1-\lambda_w \beta)}{\lambda_w}$. Union profits paid to workers therefore are $\Pi_t^U = (w_t^{fi} - w_t) N_t$.

Consumption Good Bundler

The consumption goods are bundles of domestically produced and imported final goods and are not traded across countries. Letting F_t denote the consumption good and A_t and B_t bundles of domestically and imported final goods, we assume the following aggregation technology

$$F_t = \left\{ (1-(1-n)\omega_A)^{\frac{1}{\sigma}} A_t^{\frac{\sigma-1}{\sigma}} + ((1-n)\omega_A)^{\frac{1}{\sigma}} B_t^{\frac{\sigma-1}{\sigma}} \right\}^{\frac{\sigma}{1-\sigma}}, \tag{C.32}$$

$$F_t^* = \left\{ (n\omega_B)^{\frac{1}{\sigma}} A_t^{\frac{\sigma-1}{\sigma}} + (1-n\omega_B)^{\frac{1}{\sigma}} B_t^{\frac{\sigma-1}{\sigma}} \right\}^{\frac{\sigma}{1-\sigma}}. \tag{C.33}$$

Here σ measures the terms of trade elasticity of the relative demand for domestically produced goods. $\omega_A \in [0,1]$ provides a measure for the home bias, in the sense that with $\omega_A = 1$, the Country A has no home bias. The bundles of domestically and imported final goods are defined as follows:

$$A_t = \left[\left(\frac{1}{n_A}^{\frac{1}{\sigma}} \int_0^{n_A} A_t(j)^{\frac{\epsilon-1}{\epsilon}} dj \right) \right]^{\frac{\epsilon}{\epsilon-1}}, \; B_t = \left[\left(\frac{1}{1-n_A}^{\frac{1}{\sigma}} \int_{n_A}^1 B_t(j)^{\frac{\epsilon-1}{\epsilon}} dj \right) \right]^{\frac{\epsilon}{\epsilon-1}}, \tag{C.34}$$

where $A_t(j)$ and $B_t(j)$ denote final goods produced in Home and Foreign, respectively, and ϵ measures the elasticity of substitution between final goods produced within the same country. Let $P(j)$ denote the price of a final good expressed in domestic currency. Then, letting $\mathcal{E}_t$ denote the nominal exchange rate (the price of domestic currency in terms of foreign currency) and assuming that the law of one price holds, we have

$$P_t^*(j) = \mathcal{E}_t P_t(j), \tag{C.35}$$

with $\mathcal{E}_t = 1 \; \forall t$ since both countries form a monetary union.

The optimization problem of the good bundler is to minimize expenditures subject to $F_t = C_t + I_t$, and the aggregation technologies (C.32) and (C.34). Assuming that government consumption, G_t, is a bundle that is isomorphic to consumption goods, but consists of domestically produced goods only, global demand for a generic final good produced in Country A and B are given, respectively, by

$$Y_t^d(j) = \left(\frac{P_t(j)}{P_{At}} \right)^{-\epsilon} \left\{ \left(\frac{P_{At}}{P_t} \right)^{\sigma} (1-(1-n)\omega_A)(C_t + I_t) + (1-n)\omega_B Q_t^{-\sigma}(I_t^* + C_t^*) + G_t \right\}, \tag{C.36}$$

$$Y_t^d(j)^* = \left(\frac{P_t(j)^*}{P_{Bt}^*} \right)^{-\epsilon} \left\{ \left(\frac{P_{Bt}^*}{P_t^*} \right)^{\sigma} (n\omega_A) Q_t^{\sigma}(C_t + I_t) + (1-n\omega_B)(I_t^* + C_t^*) + G_t^* \right\}, \tag{C.37}$$

where the price indices are given by

$$P_{At} = \left[\frac{1}{n} \int_0^{n_A} P_t(j)^{1-\epsilon} dj \right]^{\frac{1}{1-\epsilon}}, \; P_{Bt} = \left[\frac{1}{1-n} \int_{n_A}^1 P_t(j)^{1-\epsilon} dj \right]^{\frac{1}{1-\epsilon}} \tag{C.38}$$

and

$$P_t = [(1-(1-n)\omega_A)P_{At}^{1-\sigma} + ((1-n)\omega_A)P_{Bt}^{1-\sigma}]^{\frac{1}{1-\sigma}}, \tag{C.39}$$

$$P_t^* = [(n\omega_B)(P_{At}^*)^{1-\sigma} + (1-n\omega_B)(P_{Bt}^*)^{1-\sigma}]^{\frac{1}{1-\sigma}}. \tag{C.40}$$

The real exchange rate is given by

$$Q_t = \frac{P_t \mathcal{E}_t}{P_t^*}. \tag{C.41}$$

Final goods producers

Similar to unions, final goods producers in the home country differentiate the homogeneous home intermediate goods and set prices. They face the global demand (C.36) for each good $j \in [0, n]$ and buy the intermediate good at the national nominal price, MC_t. As we do for unions, we assume price adjustment frictions à la Calvo (1983) with indexation.

Under this assumption, the firms' managers maximize the present value of real profits given this price adjustment friction, i.e., they maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \lambda_Y^t (1-\tau_t) \left(\frac{p_{jt}(\bar{\pi})^t}{P_t} - \frac{MC_t}{P_t} \right) Y_t^d(j) \tag{C.42}$$

with a time-constant discount factor.

The corresponding first-order condition for price setting implies a domestic Phillips curve

$$\log\left(\frac{\pi_{At}}{\bar{\pi}}\right) = \beta \mathbb{E}_t \log\left(\frac{\pi_{At+1}}{\bar{\pi}}\right) + \kappa_Y \left(mc_t - \frac{1}{\mu^Y} \right) \tag{C.43}$$

where we again dropped all terms irrelevant for a first-order approximation and have $\kappa_Y = \frac{(1-\lambda_Y)(1-\lambda_Y\beta)}{\lambda_Y}$. Here, $\pi_{At} := \frac{P_{At}}{P_{At-1}}$, is the gross domestic producer price inflation rate, i.e., the gross inflation rate of domestic final goods, $mc_t := \frac{MC_t}{P_t}$ are the domestic real marginal costs, $\bar{\pi}$ is steady-state inflation, and $\frac{1}{\mu^Y} = \frac{\eta-1}{\eta}$ is the target markup. National profits paid to domestic entrepreneurs therefore are $\Pi_t^F = (1-mc_t)Y_t$.

Intermediate goods producers

Intermediate goods are produced with a constant returns to scale production function:

$$Y_t = Z_t(N_t)^{\alpha}(u_t K_t)^{(1-\alpha)} \tag{C.44}$$

where Z_t is national total factor productivity and follows an autoregressive process in logs and $u_t K_t$ is the effective capital stock taking into account utilization, u_t, i.e., the intensity with which the existing capital stock is used. Using capital with an intensity higher than normal increases depreciation of capital according to $\delta(u_t) = \delta_0 + \delta_1(u_t - 1) + \delta_2/2(u_t - 1)^2$, which, assuming $\delta_1, \delta_2 > 0$, is an increasing and convex function of utilization. Without loss of generality, capital utilization in the steady state is normalized to 1, so that δ_0 denotes the steady-state depreciation rate of capital goods.

Let mc_t be the relative price at which the intermediate good is sold to final goods producers. The intermediate goods producer maximizes profits,

$$mc_t Z_t Y_t - w_t^{fi} N_t - [r_t^F + q_t \delta(u_t)] K_t, \tag{C.45}$$

where r_t^F and q_t are the rental rate of firms and the (producer) price of capital goods, respectively. The intermediate goods producer operates in perfectly competitive national markets, such that the real wage and the user costs of capital are given by the marginal product of labor and effective capital:

$$w_t^{fi} = \alpha mc_t Z_t \left(\frac{u_t K_t}{N_t}\right)^{1-\alpha} \tag{C.46}$$

$$r_t^F + q_t \delta(u_t) = u_t (1-\alpha) mc_t Z_t \left(\frac{N_t}{u_t K_t}\right)^{\alpha} \tag{C.47}$$

We assume that utilization is decided by the owners of the capital goods, taking the aggregate national supply of capital services as given. The optimality condition for utilization is given by

$$q_t[\delta_1 + \delta_2(u_t - 1)] = (1-\alpha) mc_t Z_t \left(\frac{N_t}{u_t K_t}\right)^{\alpha}, \tag{C.48}$$

i.e., capital owners increase utilization until the marginal maintenance costs equal the marginal product of capital services.

Capital goods producers

Capital goods producers take the relative price of capital goods, q_t, as given in deciding about their output, i.e., they maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t I_t \left\{ q_t \left[1 - \frac{\phi}{2} \left(\log \frac{I_t}{I_{t-1}} \right)^2 \right] - 1 \right\}. \tag{C.49}$$

Optimality of the capital goods production requires (again dropping all terms irrelevant up to first order)

$$q_t \left[1 - \phi \log \frac{I_t}{I_{t-1}} \right] = 1 - \beta \mathbb{E}_t \left[q_{t+1} \psi \log \left(\frac{I_{t+1}}{I_t} \right) \right], \tag{C.50}$$

and each capital goods producer will adjust its production until (C.50) is fulfilled.

Since all capital goods producers within a country are symmetric, we obtain the law for motion for domestic aggregate capital as

$$K_t - (1 - \delta(u_t))K_{t-1} = \left[1 - \frac{\phi}{2}\left(\log \frac{I_t}{I_{t-1}}\right)^2\right] I_t \tag{C.51}$$

The functional form assumption implies that investment adjustment costs are minimized and equal to 0 in steady state.

C.3.3 Government Sector

The two countries form a monetary union such that they run a common monetary authority. In addition, each country runs a national fiscal authority. The monetary authority controls the nominal interest rate on liquid assets in both countries, while the national fiscal authorities issue government bonds in a union-wide bond market to finance deficits, choose both the average tax rate and the tax progressivity in their country, and make expenditures for government consumption and their national transfer system.

Monetary Union

We assume that monetary policy sets the nominal interest rate, which is the same in both countries, following a Taylor (1993)-type rule with interest rate smoothing:

$$\frac{R^b_{t+1}}{\bar{R}^b} = \left(\frac{R^b_t}{\bar{R}^b}\right)^{\rho_R} \left(\frac{n\pi_{At} + (1-n)(\pi_{Bt})}{\bar{\pi}}\right)^{(1-\rho_R)\theta_\pi} \left(n\frac{Y_t}{Y_{t-1}} + (1-n)\frac{Y^*_t}{Y^*_{t-1}}\right)^{(1-\rho_R)\theta_Y} \epsilon^R_t. \tag{C.52}$$

The coefficient $\bar{R}^b \geq 0$ determines the nominal interest rate in the steady state, Y^*_t determines output in Country B, and π_{Bt} is the producer price inflation in Country B. The coefficients $\theta_\pi, \theta_Y \geq 0$ govern the extent to which the central bank attempts to stabilize producer price inflation and the output growth in the monetary union. $\rho_R \geq 0$ captures interest rate smoothing and ϵ^R_t is an i.i.d. monetary policy shock.

Fiscal Policy

The budget constraint of the national fiscal policy reads

$$G_t + Tr_t = B_{t+1} + T_t - \frac{R^b_t}{\pi^{CPI}_t} B_t. \tag{C.53}$$

Hence, the government has expenditure for government spending, G_t, aggregate spending on its transfer system specified below, Tr_t, and repaying its debt, B_t. It finances its expenditures by issuing new debt and tax revenue, T_t. Tax revenue is

$$T_t = \tau_t(w_t N_t + \mathbb{I}_{h_{it}=0}\Pi_t^{fi} + \mathbb{I}_{h_{it}\neq 0}\Pi_t^{U}). \tag{C.54}$$

We assume that the average tax rate is a feedback function of government debt:

$$\frac{\tau_t}{\bar{\tau}} = \left(\frac{B_{t+1}}{\bar{B}}\right)^{\gamma_B^\tau}, \tag{C.55}$$

where γ_B^τ governs the speed with which debt returns to its target.

Targeted Transfer System

The targeted transfer system provides additional resources if net labor income $w_t n_t h_{it}$ falls short of some target level. For simplicity, we assume that these transfers are non-distortionary for the labor supply decision. In particular, we assume that transfers are paid to households according to the following scheme:

$$tr_{it} = \max\{0, a_1\bar{y} - a_2(1-\tau_t)w_t h_{it} n_{it}\}, \tag{C.56}$$

where $\bar{y}$ is the median income and $0 \leq a_1, a_2 \leq 1$. Thus, transfers decrease in individual income with a transfer withdrawal rate of a_2 and no transfers are paid to households whose net labor income $(1-\tau_t)w_t h_{it} n_{it} \geq \frac{a_1}{a_2}\bar{y}$. Total transfer payments of the government in Country A are then

$$Tr_t = \mathbb{E}_t tr_{it}, \tag{C.57}$$

where again, the expectation operator is the cross-sectional average.

C.3.4 Goods, bonds, capital, and labor market clearing

The national labor market in Country A clears at the competitive wage given in (C.46). A symmetric labor market clearing condition is in place in Country B. The bond markets clear whenever the following equations hold:

$$\begin{aligned} B_{t+1} &= B^d(R_t^b, r_t, q_t, \Pi_t^{fi}, \Pi_t^U, w_t, \pi_t, \tau_t, \Theta_t, \Theta_t^*, \mathbb{W}_{t+1}) - \frac{B_{Bt+1}}{Q_t} \\ &:= \mathbb{E}_t[\lambda \mathbb{B}_{a,t} + (1-\lambda)\mathbb{B}_{n,t}] - \frac{B_{Bt+1}}{Q_t}, \\ B_{t+1}^* &= B^{d,*}(R_t^b, r_t^*, q_t^*, \Pi_t^{fi,*}, \Pi_t^{U,*} w_t^*, \pi_t^{CPI,*}, \tau_t^*, \Theta_t, \Theta_t^*, \mathbb{W}_{t+1}^*) + \frac{n_A}{1-n_A} B_{Bt+1} \\ &:= \mathbb{E}_t[\lambda \mathbb{B}_{a,t}^* + (1-\lambda)\mathbb{B}_{n,t}^*] + \frac{n_A}{1-n_A} B_{Bt+1}, \\ B_{t+1}^d + B_{t+1}^{d,*} &= B_{t+1} + B_{t+1}^* \end{aligned} \tag{C.58}$$

where $\mathbb{B}_{a,t}$, $\mathbb{B}_{n,t}$ are functions of the states (b, k, h), and depend on how the households in the Country A value asset holdings in the future, $\mathbb{W}_{t+1}$, and the current set of prices (and tax rates) $(R_t^b, r_t, q_t, \Pi_t^{fi}, \Pi_t^U, w_t, \pi_t^{CPI}, \tau_t)$.[105] Future prices do not show up because we can express the value functions such that they summarize all relevant information on the expected future price paths. Expectations in the right-hand-side expression are taken w.r.t. the distributions in both countries $\Theta_t(b, k, h)$ and $\Theta_t^*(b, k, h)$. The total net amount of foreign bond holdings in Country A, B_{Bt}, is given by the aggregation over the households' budget constraint:

$$\begin{aligned}(1-\tau_t)(w_t N_t + \Pi_t^U + \Pi_t^{fi}) + (P_{At}Y_t - w_t N_t - (\Pi_t^U + \Pi_t^{fi})) + Tr_t + B_t R_t^b/\pi_t^{CPI} \\ + B_{Bt}R_t^b/(\pi_t^{CPI,*}Q_t)) = C_t + I_t + \bar{R} * BD_t + B_{t+1} + B_{Bt+1}/Q_t, \quad \text{(C.59)}\end{aligned}$$

where BD_t is the total amount of borrowing in Country A. Since both government bonds pay the same interest rate, we do not need to take track of the share of domestic vs. foreign bond holdings in each household's portfolio. Equilibrium requires the total *net* amount of bonds the household sectors in both countries demand to equal the supply of government bonds. In gross terms there are more liquid assets in circulation as some households borrow up to $\underline{B}$.

In addition, the national markets for capital have to clear. In Country A, we have:

$$\begin{aligned}K_{t+1} = K^d(R_t^b, r_t, q_t, \Pi_t^{fi}, \Pi_t^U, w_t, \pi_t^{CPI}, \tau_t, \Theta_t, \Theta_t^*, \mathbb{W}_{t+1}) \\ := \mathbb{E}_t[\lambda(\mathbb{K}_t) + (1-\lambda)(k)] \quad \text{(C.60)}\end{aligned}$$

where the first equation stems from competition in the production of capital goods, and the second equation defines the aggregate supply of funds from households in Country A - both those that trade capital, $\lambda(\mathbb{K}_t)$ and those that do not, $(1-\lambda)(k)$. Again $\mathbb{K}_t$ is a function of the current prices and continuation values. In Country B, the capital market clearing condition is symmetric.

Finally, goods market clearing requires:

$$\begin{aligned}Y_t = ((1-(1-n)\omega_A)\left(\frac{P_{At}}{P_t}\right)^{-\sigma}\left[C_t + I_t + BD_t\bar{R}\right] + (1-n_A)\omega_B Q_t^{-\sigma}\left[C_t^* + I_t^* + BD_t^*\bar{R})\right] + G_t \\ Y_t^* = n\omega_A Q_t^{\sigma}\left(\frac{P_{Bt}^*}{P_t^*}\right)^{-\sigma}\left[C_t + I_t + BD_t\bar{R}\right] + (1-n_A\omega_B)\left[C_t^* + I_t^* + BD_t^*\bar{R})\right] + G_t^*. \\ \text{(C.61)}\end{aligned}$$

C.3.5 Equilibrium

A sequential equilibrium with recursive planning in our two-country model is a sequence of policy functions $\{\mathbb{X}_{at}, \mathbb{X}_{nt}, \mathbb{B}_{at}, \mathbb{B}_{nt}, \mathbb{K}_t\}$ in Country A and $\{\mathbb{X}_{at}^*, \mathbb{X}_{nt}^*, \mathbb{B}_{at}^*, \mathbb{B}_{nt}^*, \mathbb{K}_t^*\}$ in Country

[105] The same logic applies for $\mathbb{B}_{a,t}^*$, $\mathbb{B}_{n,t}^*$ in Country B.

B, a sequence of value functions $\{V_t^a, V_t^n\}$ in Country A and $\{V_t^{a,*}, V_t^{n,*}\}$ in Country B, a sequence of prices
$\{w_t, w_t^{fi}, \Pi_t^U, \Pi_t^{fi}, q_t, r_t, R_t^b, \pi_t^{CPI}, \pi_{At}, \pi_t^W, \frac{P_{At}}{P_t}, \tau_t, Q_t\}$ in Country A and
$\{w_t^*, w_t^{fi,*}, \Pi_t^{U,*}, \Pi_t^{fi,*}, q_t^*, r_t^*, \pi_t^{CPI,*}, \pi_{Bt}, \pi_t^{W,*}, \frac{P_{Bt}^*}{P_t^*}, \tau_t^*\}$ in Country B, a sequence of of the shock ϵ_t^R, aggregate capital, labor supply, and foreign bond holdings $\{K_t, N_t, B_{Bt}\}$ in Country A and $\{K_t^*, N_t^*\}$ in Country B, distributions Θ_t in Country A and Θ_t^* in Country B over individual asset holdings and productivity, and expectations for the distribution of future prices, Γ, such that

1. Given the functionals $\mathbb{E}_t \mathbb{W}_{t+1}$ and $\mathbb{E}_t \mathbb{W}_{t+1}^*$ for the continuation value and period-t prices, policy functions $\{\mathbb{X}_{at}, \mathbb{X}_{nt}, \mathbb{B}_{at}, \mathbb{B}_{nt}, \mathbb{K}_t\}$ and $\{\mathbb{X}_{at}^*, \mathbb{X}_{nt}^*, \mathbb{B}_{at}^*, \mathbb{B}_{nt}^*, \mathbb{K}_t^*\}$ solve the households' planning problem; and given the policy functions $\{\mathbb{X}_{at}, \mathbb{X}_{nt}, \mathbb{B}_{at}, \mathbb{B}_{nt}, \mathbb{K}_t\}$ and $\{\mathbb{X}_{at}^*, \mathbb{X}_{nt}^*, \mathbb{B}_{at}^*, \mathbb{B}_{nt}^*, \mathbb{K}_t^*\}$ and prices, the value functions $\{V_t^a, V_t^n\}$ and $\{V_t^{a,*}, V_t^{n,*}\}$ are a solution to the Bellman equation.

2. Distributions of wealth and income evolve according to households' policy functions.

3. All markets clear in every period, interest rates on bonds are set according to the central bank's Taylor rule, fiscal policies are set according to the fiscal rules, and stochastic processes evolve according to their law of motion.

4. Expectations are model consistent.

We solve the model by using the perturbation method in Bayer and Luetticke (2020).

Table C.1: Calibration—Asymmetric Parameters

	Description	Country A: Italy	Country B: Germany
a_1	Transfer level	0	0.5
a_2	Transfer withdrawal rate	0	0.8
G/Y	Gov. cons. share	0.21	0.20
σ_h	STD labor inc.	0.123	0.135
β	Discount factor	0.9854	0.9823
λ	Portfolio adj. prob.	0.038	0.071
ζ	Trans. prob. from W to E	0.0007	0.001
ι	Trans prob. E to W	0.0625	0.0625
$\bar{R}$	Borrowing penalty	0.018	0.029

C.4 Calibration

We calibrate the two countries in our model to match the wealth distributions in Germany and Italy. This requires asymmetric calibration choices regarding the households. Table C.1 shows the calibration choices required for our calibration strategy which is described in Section 4.

C.4.1 Calibration of asymmetric parameters

In order to match the data, the model requires German households to be slightly less patient, asset markets (this means housing markets for most households) to be less liquid, and borrowing penalties to be higher. Yet, the mass of entrepreneurs is larger such that pure profit incomes are smaller. The level of competition (in the sense of monopolistic competition) is higher.

C.4.2 Calibration of symmetric parameters

We keep the rest of the calibration symmetric. We calibrate the parameters by matching long-run averages and using standard parameters from the literature. Table C.2 summarizes our calibration of those parameters. We calibrate to quarterly frequency.

The labor share in production, α, is 68% corresponding to a labor income share of 62%, given a markup of 10% due to an elasticity of substitution between differentiated goods of 11. The elasticity of substitution between labor varieties is also set to 11, yielding a wage markup of 10%. The parameter δ_1 that governs the cyclicality of utilization is set to 5.0. The investment adjustment cost parameter is set to 4.0. We set the Calvo parameters for price and wage adjustment probability both to 0.25. All these parameter choices are standard values in the literature.

We set relative risk aversion, ξ, to 4, following Kaplan and Violante (2014) and the Frisch elasticity, γ to 0.5 following Chetty et al. (2011). The persistence of idiosyncratic

Table C.2: Calibration—Symmetric Parameters

	Description	Value	Source/Target
Firms			
α	Share of labor	0.68	62% lab. income
η	Elast. of substitution	11	10% Price markup
η_W	Elast. of substitution	11	10% Wage markup
κ	Price adj. prob.	0.25	1 year avg. price duration
κ_W	Wage adj. prob.	0.25	1 year avg. wage duration
ϕ	Inv. adj. cost	4.0	Bayer et al. (ming)
δ_0	Depreciation rate	0.018	Wealth Gini = 0.61
δ_1	Depr. rate increase	5.0	Bayer et al. (ming)
Households			
ξ	Risk aversion	4	Kaplan and Violante (2014)
γ	Inv. Frisch elast.	2	Chetty et al. (2011)
Open economy			
σ	Trade-price elasticity	0.66	Standard value
ω	Home bias	0.66	Standard value
n	Country size	0.5	Same size
Government			
$\bar{\tau}$	Tax rate	0.3	Standard value
$\bar{R}^b$	Gross interest rate	1.00	zero interest-growth difference
ρ_R	Pers. in Taylor rule	0.75	standard value
θ_π	Reaction to Infl.	1.25	standard value
θ_Y	Reaction to Output	0	ECB mandate

Notes: Parameter values for baseline calibration. Symmetric countries.

income shocks is set to $\rho_h = 0.9815$. The stationary equilibrium real rate(-growth difference) is set to a net rate of zero.

The steady state tax level is set to 0.3. We assume that monetary policy only targets inflation, as this is the primary mandate of the ECB, and set the Taylor coefficient to 1.5 and the smoothing parameter to 0.75. The steady state inflation is zero. We assume both countries are equally large and set $n = 0.5$. The home bias parameter, ω, and the terms of trade elasticity, σ are both set to 0.66—again standard values in the literature.

C.5 The Effects of a Monetary Union with Symmetric Countries

This appendix shows the results using a symmetric calibration of our medium-sized HANK^2 model. Figure C.1 shows the IRFs after a contractionary TFP shock in Country A. Comparing the union-wide aggregates under a monetary union and with independent monetary policies reveals that our Proposition 1 holds exactly in that version of the model.

Figure C.2 shows that the same is true for Propositions 2 and 3: the bars of each wealth deciles in both countries are exactly the same size. If we do not depict absolute values,

Figure C.1: Symmetric countries: Impulse responses to TFP shock in Country H

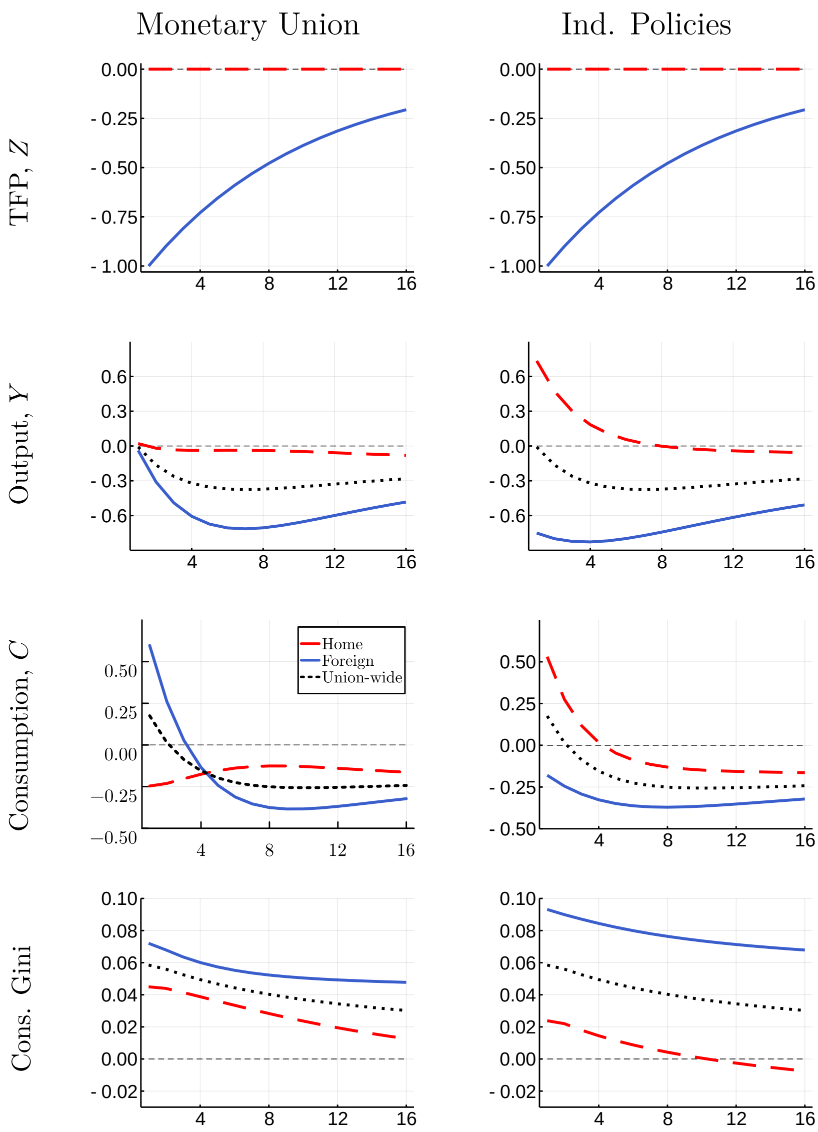

Notes: Effect of TFP shock in Country H in monetary union (left) and with independent monetary policies (right) in the model with symmetric countries. Y-axis: Percentage deviation from steady state. X-axis: Quarters.

one would also see that they always have the opposite size. Hence, the welfare effect of a monetary union on a given union-wide wealth decile is always zero (Proposition 2) implying that a monetary union only redistributes within wealth brackets across countries.

Figure C.2: Welfare impact of monetary union along the wealth distribution

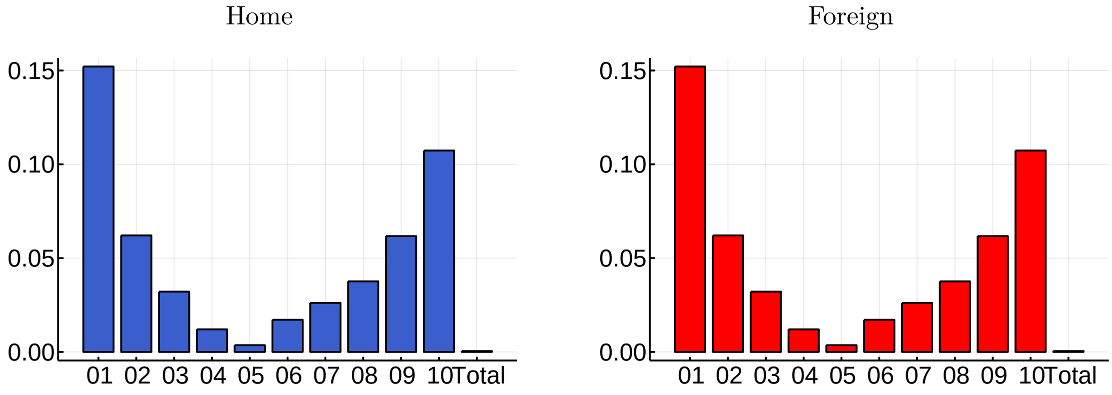

Notes: Welfare impact of monetary union measured as the difference to welfare under independent monetary policies in absolute value after TFP shocks in Country H (upper panels) and after TFP shocks in Country F (lower panel) in model with symmetric countries. Y-axis: Difference in consumption compensation (in absolute values). X-axis: Wealth percentiles.

C.6 Results for Government Spending Shock at Home

Figure C.3: Adjustment to adverse government spending shock originating in Home

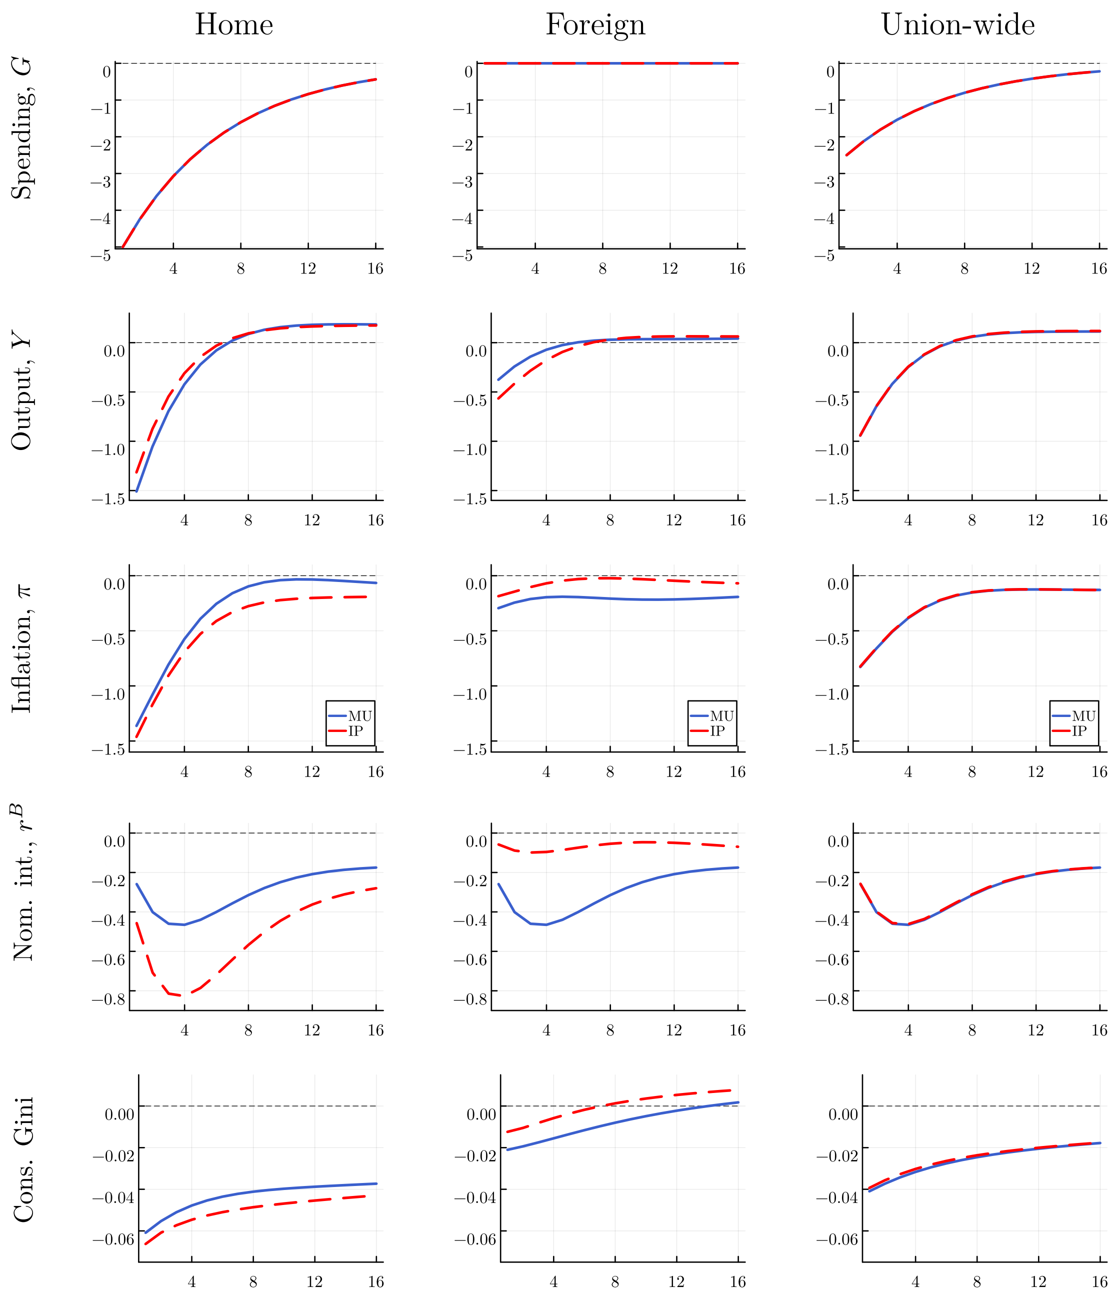

Notes: monetary union v independent monetary policies in Home (left), in Foreign (middle), and aggregate of Home and Foreign (right). Y-axis: Percentage deviation from steady state and percentage points in case of interest rates. X-axis: Quarters.

Figure C.4: How monetary union alters the welfare impact of shocks: G Shock at Home

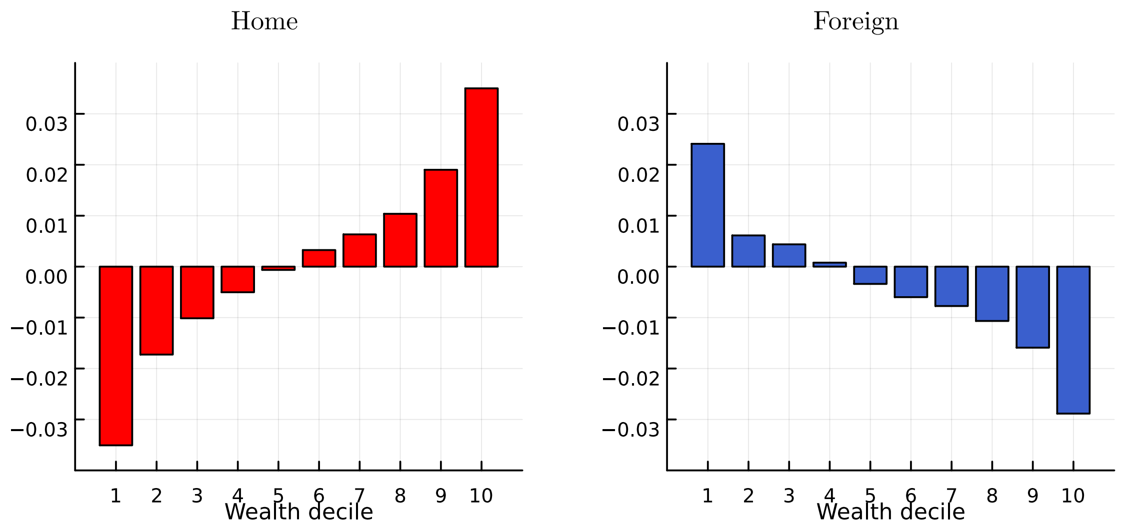

Notes: Difference of welfare impact of a government spending shock at Home between monetary union and independent monetary policies in Home (left) and Foreign (right). Y-axis: Difference in consumption compensation. X-axis: Wealth percentiles.

Appendix D

Appendix for Chapter 4

D.1 Analytical Results: Proofs and Details

D.1.1 Derivation of χ

In Section 4.3, we stated that

$$\widehat{c}_t^H = \chi \widehat{y}_t, \tag{D.1}$$

where $\chi \equiv 1 + \varphi \left(1 - \frac{\mu^D}{\lambda}\right)$ is *the* crucial statistic coming from the limited heterogeneity setup. We now show how we arrive at equation (D.1) from the H-household's budget constraint, optimality conditions and market clearing.

The labor-leisure condition of the H households is given by $(N_t^H)^{\varphi} = W_t (C_t^H)^{-\gamma}$, and similarly for the U households. As we focus on the steady state with no inequality, we have that in steady state $C = C^H = C^U$ and $N = N^U = N^H$ and market clearing and the production function imply $Y = C = N$, which we normalize to 1.

Log-linearizing the labor-leisure conditions yields $\varphi \widehat{n}_t^H = \widehat{w}_t - \gamma \widehat{c}_t^H$ and $\varphi \widehat{n}_t^U = \widehat{w}_t - \gamma \widehat{c}_t^U$. Since both households work for the same wage, we obtain

$$\varphi \widehat{n}_t^H + \gamma \widehat{c}_t^H = \varphi \widehat{n}_t^U + \gamma \widehat{c}_t^U \tag{D.2}$$

Log-linearizing the market clearing conditions yields $\widehat{n}_t = \lambda \widehat{n}_t^H + (1-\lambda) \widehat{n}_t^U$ and $\widehat{c}_t = \lambda \widehat{c}_t^H + (1-\lambda) \widehat{c}_t^U$, which can be re-arranged as (using $\widehat{y}_t = \widehat{c}_t = \widehat{n}_t$)

$$\begin{aligned} \widehat{n}_t^U &= \frac{1}{1-\lambda} \left(\widehat{y}_t - \lambda \widehat{n}_t^H \right) \\ \widehat{c}_t^U &= \frac{1}{1-\lambda} \left(\widehat{y}_t - \lambda \widehat{c}_t^H \right). \end{aligned}$$

Replacing $\widehat{n}_t^U$ and $\widehat{c}_t^U$ in equation (D.2) then gives

$$\varphi \widehat{n}_t^H + \gamma \widehat{c}_t^H = (\varphi + \gamma) \widehat{y}_t. \tag{D.3}$$

The budget constraint of H households (accounting for the fact that bond holdings are zero in equilibrium) is given by $C_t^H = W_t N_t^H + \frac{\mu^D}{\lambda} D_t$. In log-linearized terms, we get

$$\widehat{c}_t^H = \widehat{w}_t + \widehat{n}_t^H + \frac{\mu^D}{\lambda}\widehat{d}_t, \tag{D.4}$$

and using that $\widehat{w}_t = -\widehat{d}_t = \varphi \widehat{n}_t^H + \gamma \widehat{c}_t^H$, we get

$$\widehat{c}_t^H = \left(\varphi \widehat{n}_t^H + \gamma \widehat{c}_t^H\right)\left(1 - \frac{\mu^D}{\lambda}\right) + \widehat{n}_t^H. \tag{D.5}$$

Using (D.3) to solve for $\widehat{n}_t^H$ and plugging it into (D.5) yields

$$\widehat{c}_t^H = \widehat{c}_t^H \gamma \left(1 - \frac{\mu^D}{\lambda}\right) + \chi\left(\frac{\varphi + \gamma}{\varphi}\widehat{y}_t - \frac{\gamma}{\varphi}\widehat{c}_t^H\right).$$

Grouping terms, we obtain

$$\widehat{c}_t^H = \chi \widehat{y}_t,$$

with $\chi \equiv 1 + \varphi\left(1 - \frac{\mu^D}{\lambda}\right)$, as stated above.

D.1.2 Proof of Proposition 5.

When linearizing the model around the steady state, our bounded rationality assumptions imply

$$\mathbb{E}_t^{BR}\left[\widehat{x}_{t+1}\right] = \bar{m}\mathbb{E}_t\left[\widehat{x}_{t+1}\right]. \tag{D.6}$$

Combining equations (4.12) and (4.14) with (D.6), we have

$$\begin{aligned}
\mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^H\right] &= \bar{m}\mathbb{E}_t\left[\widehat{c}_{t+1}^H\right] = \bar{m}\chi\mathbb{E}_t\left[\widehat{y}_{t+1}\right] \\
\mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^U\right] &= \bar{m}\mathbb{E}_t\left[\widehat{c}_{t+1}^U\right] = \bar{m}\frac{1-\lambda\chi}{1-\lambda}\mathbb{E}_t\left[\widehat{y}_{t+1}\right].
\end{aligned}$$

Plugging these two equations as well as equation (4.14) into the Euler equation of unconstrained households (4.16) yields

$$\frac{1-\lambda\chi}{1-\lambda}\widehat{y}_t = s\bar{m}\frac{1-\lambda\chi}{1-\lambda}\mathbb{E}_t\left[\widehat{y}_{t+1}\right] + (1-s)\bar{m}\chi\mathbb{E}_t\left[\widehat{y}_{t+1}\right] - \frac{1}{\gamma}\left(\widehat{i}_t - \mathbb{E}_t\pi_{t+1}\right).$$

(Note, that this is exactly the same expression as in section 4.3.5 but with χ instead of ζ.) Combining the $\mathbb{E}_t\left[\widehat{y}_{t+1}\right]$ terms and dividing by $\frac{1-\lambda\chi}{1-\lambda}$ yields the following coefficient in front of $\mathbb{E}_t\left[\widehat{y}_{t+1}\right]$:

$$\begin{aligned}\psi_f &\equiv \bar{m}\left[s+(1-s)\chi\frac{1-\lambda}{1-\lambda\chi}\right]=\bar{m}\left[1-1+s+(1-s)\chi\frac{1-\lambda}{1-\lambda\chi}\right]\\ &=\bar{m}\left[1-\frac{1-\lambda\chi}{1-\lambda\chi}+s+(1-s)\chi\frac{1-\lambda}{1-\lambda\chi}\right]=\bar{m}\left[1-\frac{1-\lambda\chi}{1-\lambda\chi}+\frac{(1-\lambda\chi)s}{1-\lambda\chi}+(1-s)\chi\frac{1-\lambda}{1-\lambda\chi}\right]\\ &=\bar{m}\left[1+(\chi-1)\frac{1-s}{1-\lambda\chi}\right].\end{aligned}$$

Defining $\psi_c \equiv \frac{1-\lambda}{1-\lambda\chi}$ yields the behavioral HANK IS equation in Proposition 5:

$$\widehat{y}_t=\psi_f\mathbb{E}_t\widehat{y}_{t+1}-\psi_c\frac{1}{\gamma}\left(\widehat{i}_t-\mathbb{E}_t\pi_{t+1}\right).$$

D.1.3 Proof of Proposition 6.

We prove here the more general case where forward guidance is implemented as changes in the nominal rather than the real rate and where the supply side is captured by the Phillips Curve (4.17). The case with real rate changes is a special case of the nominal rate case and can be captured by setting $\kappa = 0$.

The first part of Proposition 6 follows from the fact that amplification is obtained when

$$\psi_c=\frac{1-\lambda}{1-\lambda\chi}>1,$$

which requires $\chi > 1$, given that we assume throughout $\chi\lambda < 1$.

For the second part, recall how we define the forward guidance experiment (following Bilbiie (2021)). We assume a Taylor coefficient of 0, i.e., $\phi = 0$, such that the nominal interest rate is given by $\widehat{i}_t = \varepsilon_t^{MP}$. Replacing inflation using the Phillips curve (4.17), i.e., $\pi_t = \kappa\widehat{y}_t$, we can re-write the behavioral HANK IS equation from Proposition 5 as

$$\begin{aligned}\widehat{y}_t &= \psi_f\mathbb{E}_t\widehat{y}_{t+1}-\psi_c\frac{1}{\gamma}\left(\varepsilon_t^{MP}-\kappa\mathbb{E}_t\widehat{y}_{t+1}\right)\\ &=\left(\psi_f+\psi_c\frac{1}{\gamma}\kappa\right)\mathbb{E}_t\widehat{y}_{t+1}-\psi_c\frac{1}{\gamma}\varepsilon_t^{MP}\end{aligned}$$

The forward guidance puzzle is ruled out if and only if

$$\left(\psi_f+\psi_c\frac{1}{\gamma}\kappa\right)<1,$$

which is the same as:

$$\bar{m}\delta+\frac{1}{\gamma}\frac{1-\lambda}{1-\lambda\chi}\kappa<1.$$

D.1.4 Proof of Proposition 7.

Replacing $\widehat{i}_t$ by $\phi\pi_t = \phi\kappa\widehat{y}_t$ and $\mathbb{E}_t\pi_{t+1} = \kappa\mathbb{E}_t\widehat{y}_{t+1}$ (which follows from the Taylor rule and the static Phillips Curve) in the IS equation (4.18), we get

$$\widehat{y}_t = \psi_f\mathbb{E}_t\widehat{y}_{t+1} - \psi_c\frac{1}{\gamma}\left(\phi\kappa\widehat{y}_t - \kappa\mathbb{E}_t\widehat{y}_{t+1}\right),$$

which can be re-written as

$$\widehat{y}_t\left(1 + \psi_c\frac{1}{\gamma}\phi\kappa\right) = \mathbb{E}_t\widehat{y}_{t+1}\left(\psi_f + \psi_c\frac{1}{\gamma}\kappa\right).$$

Dividing by $\left(1 + \psi_c\frac{1}{\gamma}\phi\kappa\right)$ and plugging in for ψ_f and ψ_c yields

$$\widehat{y}_t = \frac{\bar{m}\delta + \frac{(1-\lambda)\kappa}{\gamma(1-\lambda\chi)}}{1 + \kappa\phi\frac{1}{\gamma}\frac{(1-\lambda)}{1-\lambda\chi}}\mathbb{E}_t\widehat{y}_{t+1}.$$

To obtain determinacy, the term in front of $\mathbb{E}_t\widehat{y}_{t+1}$ has to be smaller than 1. Solving this for ϕ yields

$$\phi > \phi^* = 1 + \frac{\bar{m}\delta - 1}{\frac{\kappa}{\gamma}\frac{1-\lambda}{1-\lambda\chi}}, \tag{D.7}$$

which is the condition in Proposition 7. This illustrates how bounded rationality raises the likelihood that the Taylor principle ($\phi^* = 1$) is sufficient for determinacy, as the Taylor principle can only hold if $\bar{m}\delta \leq 1$. In the rational model, this boils down to $\delta \leq 1$. However, the Taylor principle can be sufficient under bounded rationality, i.e., $\bar{m} < 1$, even when $\delta > 1$, thus, even when allowing for amplification. Note that we could also express condition (D.7) as

$$\phi > \phi^* = 1 + \frac{\psi_f - 1}{\frac{\kappa}{\gamma}\psi_c}.$$

Generalizations of Proposition 7. Proposition 7 can easily be extended to allow for Taylor rules of the form

$$\widehat{i}_t = \phi_\pi\pi_t + \phi_y\widehat{y}_t$$

and in which the behavioral agents do not have rational expectations about the real interest rate but rather perceive the real interest rate to be equal to

$$\widehat{r}_t^{BR} \equiv \widehat{i}_t - \bar{m}^r\mathbb{E}_t\pi_{t+1},$$

where $\bar{m}^r$ can be equal to $\bar{m}$ or can potentially differ from it (if it equals 1, we are back to the case in which the behavioral agent is rational with respect to real interest rates).

Combining the static Phillips Curve with the generalized Taylor rule and the behavioral

HANK IS equation, it follows that

$$\widehat{y}_t = \frac{\psi_f + \frac{\kappa}{\gamma}\psi_c \bar{m}^r}{1 + \frac{\psi_c}{\gamma}\left(\kappa\phi_\pi + \phi_y\right)} \mathbb{E}_t \widehat{y}_{t+1}. \tag{D.8}$$

From equation (D.8), it follows that we need

$$\phi_\pi > \bar{m}^r - \phi_y + \frac{\psi_f - 1}{\psi_c \frac{\kappa}{\gamma}} = \bar{m}^r - \phi_y + \frac{\bar{m}\delta - 1}{\frac{1-\lambda}{1-\chi\lambda}\frac{\kappa}{\gamma}} \tag{D.9}$$

for the model to feature a determinate, locally unique equilibrium. Condition (D.9) shows that both, $\bar{m}^r < 1$ and $\phi_y > 0$, weaken the condition in Proposition 7. Put differently, bounded rationality with respect to the real rate or a Taylor rule that responds to changes in output, both relax the condition on ϕ_π to yield determinacy.

D.1.5 Derivation of Lemma 2

Let us first state a few auxiliary results that will prove helpful later. First, in log-linearized terms, the stochastic discount factor is given by

$$\frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}^U_{t,t+1} = \widehat{c}^U_t - s\bar{m}\mathbb{E}_t\widehat{c}^U_{t+1} - (1-s)\bar{m}\mathbb{E}_t\widehat{c}^H_{t+1}$$

and for i periods ahead:

$$\frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}^U_{t,t+i} = \widehat{c}^U_t - s\bar{m}^i\mathbb{E}_t\widehat{c}^U_{t+i} - (1-s)\bar{m}^i\mathbb{E}_t\widehat{c}^H_{t+i}.$$

Furthermore, we have:

$$\begin{aligned}\frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}^U_{t+1,t+2} &= \mathbb{E}_t^{BR}\left[\widehat{c}^U_{t+1} - s\widehat{c}^U_{t+2} - (1-s)\widehat{c}^H_{t+2}\right] \\ &= \bar{m}\mathbb{E}_t\widehat{c}^U_{t+1} - s\bar{m}^2\mathbb{E}_t\widehat{c}^U_{t+2} - (1-s)\bar{m}^2\mathbb{E}_t\widehat{c}^H_{t+2}\end{aligned}$$

and the stochastic discount factor has the property

$$\mathbb{E}_t^{BR}\left[\widehat{q}^U_{t,t+i}\right] = \mathbb{E}_t^{BR}\left[\widehat{q}^U_{t,t+1} + \widehat{q}^U_{t+1,t+2} + ... + \widehat{q}^U_{t+i-1,t+i}\right].$$

Using these results, $\mathbb{E}_t^{BR}\left[\widehat{q}^U_{t,t+i}\right]$ can be written as

$$\begin{aligned}\frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}^U_{t,t+i} = \widehat{c}^U_t &+ (1-s)\bar{m}\mathbb{E}_t\left[\widehat{c}^U_{t+1} - \widehat{c}^H_{t+1}\right] + (1-s)\bar{m}^2\mathbb{E}_t\left[\widehat{c}^U_{t+2} - \widehat{c}^H_{t+2}\right] + ...+ \\ &+ (1-s)\bar{m}^i\mathbb{E}_t\left[\widehat{c}^U_{t+i} - \widehat{c}^H_{t+1}\right] - \bar{m}^i\mathbb{E}_t\widehat{c}^U_{t+i},\end{aligned}$$

or put differently

$$\frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}^U_{t,t+i} + \bar{m}^i\mathbb{E}_t\widehat{c}^U_{t+i} = \widehat{c}^U_t + (1-s)\mathbb{E}_t\sum_{k=1}^{i}\bar{m}^k\left(\widehat{c}^U_{t+k} - \widehat{c}^H_{t+k}\right). \tag{D.10}$$

The (linearized) budget constraint can be written as

$$\begin{aligned}
\mathbb{E}_t^{BR}\sum_{i=0}^{\infty}\beta^i\left(\frac{1}{\gamma}\widehat{q}^U_{t,t+i} + \widehat{c}^U_{t+i}\right) &= \mathbb{E}_t^{BR}\sum_{i=0}^{\infty}\beta^i\left(\frac{1}{\gamma}\widehat{q}^U_{t,t+i} + \widehat{y}^U_{t+i}\right)\\
\Leftrightarrow \mathbb{E}_t^{BR}\sum_{i=0}^{\infty}\beta^i\left(\frac{1}{\gamma}\widehat{q}^U_{t,t+i}\right) + \mathbb{E}_t\sum_{i=0}^{\infty}(\beta\bar{m})^i\,\widehat{c}^U_{t+i} &= \mathbb{E}_t^{BR}\sum_{i=0}^{\infty}\beta^i\left(\frac{1}{\gamma}\widehat{q}^U_{t,t+i}\right) + \mathbb{E}_t\sum_{i=0}^{\infty}(\beta\bar{m})^i\,\widehat{y}^U_{t+i}.
\end{aligned}$$

Now, focus on the left-hand side and notice that the sum $\mathbb{E}_t\sum_{i=0}^{\infty}(\beta\bar{m})^i\,\widehat{c}^U_{t+i}$ cancels with the $\bar{m}^i\mathbb{E}_t\widehat{c}^U_{t+i}$ terms in equation (D.10) when summing them up. The left-hand side of the budget constraint can thus be written as

$$\begin{aligned}
&\mathbb{E}_t\sum_{i=0}^{\infty}\beta^i\imath\left(\widehat{c}^U_t + (1-s)\sum_{k=1}^{i}\bar{m}^k\left(\widehat{c}^U_{t+k} - \widehat{c}^H_{t+k}\right)\right)\\
&= \frac{1}{1-\beta}\widehat{c}^U_t + (1-s)\mathbb{E}_t\sum_{i=0}^{\infty}\beta^i\sum_{k=1}^{i}\bar{m}^k\left(\widehat{c}^U_{t+k} - \widehat{c}^H_{t+k}\right)\\
&= \frac{1}{1-\beta}\widehat{c}^U_t + \frac{1-s}{1-\beta}\mathbb{E}_t\sum_{i=1}^{\infty}(\beta\bar{m})^i\left(\widehat{c}^U_{t+i} - \widehat{c}^H_{t+i}\right).
\end{aligned}$$

Note, from the Euler equation of the unconstrained households, we obtain the real interest rate

$$-\frac{1}{\gamma}\widehat{r}_t = \widehat{c}^U_t - s\mathbb{E}_t^{BR}\widehat{c}^U_{t+1} - (1-s)\mathbb{E}_t^{BR}\widehat{c}^H_{t+1} = \frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}^U_{t,t+1},$$

and similarly,

$$-\frac{1}{\gamma}\bar{m}^i\mathbb{E}_t\widehat{r}_{t+i} = \frac{1}{\gamma}\mathbb{E}_t^{BR}\widehat{q}^U_{t+i,t+i+1},$$

where $\widehat{r}_t$ is the (linearized) real interest rate.

Combining these results, we see that

$$\mathbb{E}_t^{BR}\sum_{i=0}^{\infty}\beta^i\frac{1}{\gamma}\widehat{q}^U_{t,t+i} = -\frac{1}{1-\beta}\frac{1}{\gamma}\beta\mathbb{E}_t\sum_{i=0}^{\infty}(\beta\bar{m})^i\,\widehat{r}_{t+i}.$$

Plugging this into the right-hand side of the budget constraint and multiplying both sides by $1-\beta$ yields

$$\begin{aligned}\widehat{c}_t^U =& -\frac{1}{\gamma}\beta\widehat{r}_t + (1-\beta)\widehat{y}_t^U - (1-s)\mathbb{E}_t\sum_{i=1}^{\infty}(\beta\bar{m})^i\left(\widehat{c}_{t+i}^U - \widehat{c}_{t+i}^H\right)\\ &-\frac{1}{\gamma}\beta\mathbb{E}_t\sum_{i=1}^{\infty}(\beta\bar{m})^i\,\widehat{r}_{t+i} + (1-\beta)\mathbb{E}_t\sum_{i=1}^{\infty}(\beta\bar{m})^i\,\widehat{y}_{t+i}^U,\end{aligned}$$

or written recursively

$$\widehat{c}_t^U = -\frac{1}{\gamma}\beta\widehat{r}_t + (1-\beta)\widehat{y}_t^U + \beta\bar{m}s\mathbb{E}_t\widehat{c}_{t+1}^U + \beta\bar{m}(1-s)\mathbb{E}_t\widehat{c}_{t+1}^H.$$

Now, aggregating, i.e., multiplying the expression for $\widehat{c}_t^U$ by $(1-\lambda)$, adding $\lambda\widehat{c}_t^H$ and using $\widehat{c}_t^H = \chi\widehat{y}_t$ as well as $\widehat{y}_t^U = \frac{1-\lambda\chi}{1-\lambda}\widehat{y}_t$, yields the consumption function

$$\widehat{c}_t = [1-\beta(1-\lambda\chi)]\,\widehat{y}_t - \frac{(1-\lambda)\beta}{\gamma}\widehat{r}_t + \beta\bar{m}\delta(1-\lambda\chi)\mathbb{E}_t\widehat{c}_{t+1}, \tag{D.11}$$

as stated in the main text.

To obtain the share of indirect effects, note that the model does not feature any endogenous state variables and hence, endogenous variables inherit the persistence of the exogenous variables, ρ. Thus, $\mathbb{E}_t\widehat{c}_{t+1} = \rho\widehat{c}_t$. Plugging this into the consumption function (D.11), we get

$$\widehat{c}_t = \frac{1-\beta(1-\lambda\chi)}{1-\beta\bar{m}\delta\rho(1-\lambda\chi)}\widehat{y}_t - \frac{(1-\lambda)\beta}{\gamma(1-\beta\bar{m}\delta\rho(1-\lambda\chi))}\widehat{r}_t.$$

The term in front of $\widehat{y}_t$ is the share of indirect general equilibrium effects.

D.1.6 Derivation of the IS Equation with TFP Shocks

The production function is now given by

$$Y_t = A_t N_t, \tag{D.12}$$

where A_t denotes TFP and we assume that $A = 1$ in steady state. In log deviations, the production function reads

$$y_t = a_t + n_t. \tag{D.13}$$

We still have $Y_t = C_t$. Profits are given by

$$D_t = Y_t - W_t N_t = N_t(A_t - W_t), \tag{D.14}$$

or in log deviations

$$d_t = y_t - w_t - n_t = y_t - w_t - y_t + a_t = a_t - w_t. \tag{D.15}$$

Following the steps from D.1.2, we get

$$w_t = \varphi n_t^H + \gamma c_t^H = \varphi n_t^U + \gamma c_t^U \tag{D.16}$$

$$n_t^U = \frac{1}{1-\lambda}\left[y_t - a_t - \lambda n_t^H\right] \tag{D.17}$$

$$c_t^U = \frac{1}{1-\lambda}\left[y_t - \lambda c_t^H\right]. \tag{D.18}$$

Taking them together, we arrive at the following:

$$\varphi n_t^H + \gamma c_t^H = \frac{\varphi}{1-\lambda}\left[y_t - a_t - \lambda n_t^H\right] + \frac{\gamma}{1-\lambda}\left[y_t - \lambda c_t^H\right]. \tag{D.19}$$

Solving this for n_t^H yields

$$n_t^H = \frac{\varphi+\gamma}{\varphi} y_t - a_t - \frac{\gamma}{\varphi} c_t^H. \tag{D.20}$$

The budget constraint of the H household is given by

$$\widehat{c}_t^H = \widehat{w}_t + \widehat{n}_t^H + \frac{\mu^D}{\lambda}\widehat{d}_t, \tag{D.21}$$

which can be rewritten as

$$c_t^H = \chi n_t^H + \gamma\left(1 - \frac{\mu^D}{\lambda}\right) c_t^H + \frac{\mu^D}{\lambda} a_t. \tag{D.22}$$

Taken together, this yields

$$c_t^H = \chi y_t - (\chi - 1)\frac{1+\varphi}{\varphi+\gamma} a_t, \tag{D.23}$$

and hence

$$c_t^U = \frac{1-\lambda\chi}{1-\lambda} y_t + \frac{\lambda(\chi-1)}{1-\lambda}\frac{1+\varphi}{\varphi+\gamma} a_t. \tag{D.24}$$

Inequality is given by

$$c_t^U - c_t^H = \frac{1-\chi}{1-\lambda} y_t + \frac{1-\chi}{1-\lambda}\frac{1+\varphi}{\varphi+\gamma} a_t. \tag{D.25}$$

Plugging these two last expressions into the Euler equation, we arrive at our aggregate IS equation:

$$y_t = \psi_f E_t y_{t+1} - \frac{1}{\gamma}\psi_c r_t - (\chi-1)\frac{1+\varphi}{\varphi+\gamma}\left[\lambda + \bar{m}\rho_a(1-s-\lambda)\right] a_t. \tag{D.26}$$

Thus, as long as $\chi > 1$, TFP shocks enter the IS equation with a negative sign. A larger χ makes it more negative as does a lower $\bar{m}$, as long as $1 - s - \lambda < 0$. A sufficient condition for this is $1 < s + h$, where h is the probability of remaining H when currently H. If it is more likely to stay in the current state than switching to the other one, i.e., $s > 0.5$ and $h > 0.5$, this condition is satisfied. The cross-derivative with respect to χ and $\bar{m}$ is positive, indicating that cognitive discounting and countercyclical income risk reinforce each other.

By iterating forward equation (D.26), we obtain

$$y_t = -\frac{\psi_c}{\gamma}\sum_{s=0}^{\infty}\psi_f^s r_{t+s} - \frac{\psi_a}{1-\rho_a\psi_f}a_t, \tag{D.27}$$

where

$$\psi_a \equiv (\chi - 1)\frac{1+\varphi}{\varphi+\gamma}\left[\lambda + \bar{m}\rho_a(1-s-\lambda)\right]. \tag{D.28}$$

Derivative w.r.t. χ:

$$\frac{\delta y_t}{\delta\chi} = -\frac{\delta(\psi_c)/\delta\chi}{\gamma}\sum_{s=0}^{\infty}s(\psi_f)^{s-1}\frac{\delta\psi_f}{\delta\chi}r_{t+s} - \frac{(1-\rho_a\psi_f)*\delta\psi_a/\delta\chi - \psi_a(-\rho_a)\delta_f/\delta\chi}{(1-\rho_a\psi_f)^2}a_t. \tag{D.29}$$

The term in front of r_{t+s} is negative and the term in front of a_t is positive (given that $\chi > 1$ and $\psi_f < 1$).

Derivative w.r.t. $\bar{m}$:

$$\frac{\delta y_t}{\delta\bar{m}} = -\frac{\psi_c}{\gamma}\sum_{s=0}^{\infty}s(\psi_f)^{s-1}r_{t+s}\frac{\delta\psi_f}{\delta\bar{m}} - \frac{(1-\rho_a\psi_f)*\delta\psi_a/\delta\bar{m} - \psi_a(-\rho_a)\delta_f/\delta\bar{m}}{(1-\rho_a\psi_f)^2}a_t. \tag{D.30}$$

first term again negative, second term positive given $\chi > 1$.

Taylor with static PC ($\pi_t = \kappa x_t$):

$$y_t = \frac{\kappa\frac{\psi_c}{\gamma}\frac{1+\varphi}{\varphi+\gamma}(\phi_\pi - \rho_a) - \psi_a}{1+\frac{\psi_c}{\gamma}\kappa\phi_\pi}\frac{1}{1-\frac{\psi_f+\frac{\psi_c}{\gamma}\kappa}{1+\frac{\psi_c}{\gamma}\kappa\phi_\pi}\rho_a}a_t \tag{D.31}$$

$$= \frac{\kappa\frac{\psi_c}{\gamma}\frac{1+\varphi}{\varphi+\gamma}(\phi_\pi - \rho_a) - \psi_a}{1-\rho_a\psi_f + \frac{\psi_c}{\gamma}\kappa(\phi_\pi - \rho_a)}a_t \tag{D.32}$$

Figure D.1: The effects of a negative productivity shock on output

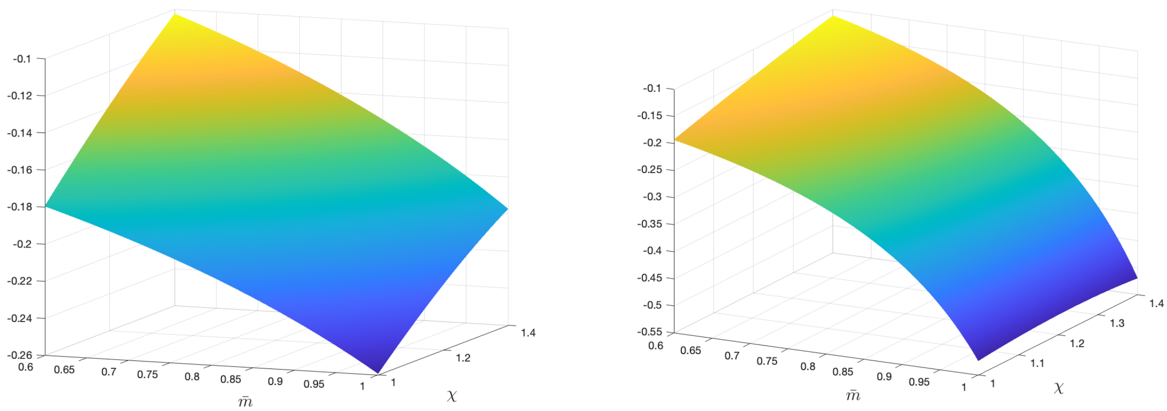

Note: This figure shows ... left: $\rho_a = 0.6$, right: $\rho_a = 0.9$. Rest of calibration: $\gamma = 1, \varphi = 1, \phi_\pi = 2$, $s = 0.8^{(}0.25), \lambda = 0.33, \kappa = 0.1$.

With sticky wages, we get (have to check this again)

$$y_t = \psi_f E_t y_{t+1} - \psi_c \frac{1}{\gamma} r_t + \frac{\chi}{1-\lambda\chi} \left[s\bar{m}\lambda\rho_a - (1-s)\bar{m}\rho_a(1-\lambda) - \lambda \right] a_t. \tag{D.33}$$

We can re-write the term in front of a_t as $\frac{\chi}{1-\lambda\chi}\left[\bar{m}\rho_a(s+\lambda-1)-\lambda\right]$ (which is identical to the term in brackets for the sticky price case.

D.2 Calibrating $\bar{m}$

In most of our analysis, we set the cognitive discounting parameter $\bar{m}$ to 0.85, as in Gabaix (2020). One way at arriving at this value is by matching estimated IS equations. Fuhrer and Rudebusch (2004), for example, estimate an IS equation and find that the coefficient in front of $\mathbb{E}_t \widehat{y}_{t+1}$ (what we call ψ_f) is approximately 0.65, which together with $\delta > 1$, would imply a $\bar{m}$ much lower than 0.85 and especially our determinacy results would be even stronger under such a calibration.

Another way to calibrate $\bar{m}$ (as pointed out in Gabaix (2020)) is to interpret the estimates in Coibion and Gorodnichenko (2015) through the "cognitive-discounting lens". They regress forecast errors on forecast revisions

$$x_{t+h} - F_t x_{t+h} = c + b^{CG}\left(F_t x_{t+h} - F_{t-1}x_{t+h}\right) + u_t,$$

where $F_t x_{t+h}$ denotes the forecast at time t of variable x, h periods ahead. Focusing on inflation, they find that $b^{CG} > 0$ in consensus forecasts, pointing to *underreaction* (similar results are, for example, found in Angeletos et al. (2021) and Adam et al. (2022) for other variables).

In the linearized model, the law of motion of x is $x_{t+1} = \Gamma\left(x_t + \varepsilon_{t+1}\right)$ whereas the behavioral agents perceive it to be $x_{t+1} = \bar{m}\Gamma\left(x_t + \varepsilon_{t+1}\right)$. It follows that $F_t x_{t+h} = \left(\bar{m}\Gamma\right)^h x_t$ and thus, forecast revisions are equal to

$$\begin{aligned} F_t x_{t+h} - F_{t-1}x_{t+h} &= \left(\bar{m}\Gamma\right)^h x_t - \left(\bar{m}\Gamma\right)^{h+1} x_{t-1} \\ &= \left(\bar{m}\Gamma\right)^h \Gamma(1-\bar{m})x_{t-1} + \left(\bar{m}\Gamma\right)^h \varepsilon_t. \end{aligned}$$

The forecast error is given by

$$x_{t+h} - F_t x_{t+h} = \Gamma^h(1-\bar{m}^h)\Gamma x_{t-1} + \Gamma^h(1-\bar{m}^h)\varepsilon_t + \sum_{j=0}^{h-1}\Gamma^j \varepsilon_{t+h-j},$$

where $\sum_{j=0}^{h-1}\Gamma^j \varepsilon_{t+h-j}$ is the rational expectations forecast error. Gabaix (2020) shows that b^{CG} is bounded below $b^{CG} \geq \frac{1-\bar{m}^h}{\bar{m}^h}$, showing that $\bar{m} < 1$ yields $b^{CG} > 0$, as found empirically.

When replacing the weak inequality with an equality, we get

$$\bar{m}^h = \frac{1}{1 + b^{CG}}.$$

Most recently, Angeletos et al. (2021) estimate b^{CG} (focusing on a horizon $h = 3$) to lie between $b^{CG} \in [0.74, 0.81]$ for unemployment forecasts and $b^{CG} \in [0.3, 1.53]$ for inflation, depending on the considered period (see their Table 1). These estimates imply $\bar{m} \in [0.82, 0.83]$ for unemployment and $\bar{m} \in [0.73, 0.92]$ for inflation, and are thus close to our preferred value of 0.85. Note, however, that these estimates pertain to professional forecasters and should therefore be seen as upper bounds on $\bar{m}$. As outlined in Section 4.2, we estimate these regressions for households to obtain more direct evidence on $\bar{m}$ for households (of different income groups). The following subsection discusses the data, the empirical strategy and the findings we obtain in more detail.

D.2.1 Estimating $\bar{m}$ for different Household Groups

To test for heterogeneity in the degree of cognitive discounting, we follow Coibion and Gorodnichenko (2015) and regress forecast errors on forecast revisions as follows

$$x_{t+4} - \mathbb{E}_t^{e,BR} x_{t+4} = c^e + b^{e,CG} \left(\mathbb{E}_t^{e,BR} x_{t+4} - \mathbb{E}_{t-1}^{e,BR} x_{t+4} \right) + \epsilon_t^e, \tag{D.34}$$

to estimate $b^{e,CG}$ for different groups of households, indexed by e. As shown above, $b^{e,CG} > 0$ is consistent with underreaction and the corresponding cognitive discounting parameter is approximately given by (we calibrate the model at quarterly frequency whereas the data is about 1-year-ahead expectations, thus, the adjustment in the exponent)

$$\bar{m}^e = \left(\frac{1}{1 + b^{e,CG}} \right)^{1/4}. \tag{D.35}$$

Ideally, we would use actual data and expectations data about future marginal utilities of consumption where changes in these variables only being driven by aggregate shocks. However, that data is not available. Instead, we focus on expectations about future unemployment (and inflation) where it seems reasonable to assume that they are only driven by aggregate shocks and that they matter for household's (actual and expected) marginal consumption utility. The Survey of Consumers from the University of Michigan provides 1-year ahead unemployment expectations and we use the unemployment rate from the FRED database as our measure of actual unemployment. We split the households into three groups based on their income. The bottom and top income groups each contain the 25% households with the lowest and highest income, respectively, and the remaining 50% are assigned to the middle income group.

The Michigan Survey asks households whether they expect unemployment to increase, decrease or to remain about the same over the next twelve months. We follow Carlson and

Parkin (1975), Mankiw (2000) and Bhandari et al. (2019) to translate these categorical unemployment expectation into numerical expectations.

Focus on group $e \in \{L, M, H\}$ and let $q_t^{e,D}$, $q_t^{e,S}$ and $q_t^{e,U}$ denote the shares within income group e reported at time t that think unemployment will go down, stay roughly the same, or go up over the next year, respectively. We assume that these shares are drawn from a cross-sectional distribution of responses that are normally distributed according to $\mathcal{N}\left(\mu_t^e, (\sigma_t^e)^2\right)$ and a threshold a such that when a household expects unemployment to remain within the range $[-a, a]$ over the next year, she responds that unemployment will remaine "about the same". We thus have

$$q_t^{e,D} = \Phi\left(\frac{-a-\mu_t^e}{\sigma_t^e}\right) \qquad q_t^{e,U} = 1 - \Phi\left(\frac{a-\mu_t^e}{\sigma_t^e}\right),$$

which after some rearranging yields

$$\begin{aligned} \sigma_t^e &= \frac{2a}{\Phi^{-1}\left(1-q_t^{e,U}\right) - \Phi^{-1}\left(q_t^{e,D}\right)} \\ \mu_t^e &= a - \sigma_t^e \Phi^{-1}\left(1-q_t^{e,U}\right). \end{aligned}$$

This leaves us with one degree of freedom, namely a. We make two assumptions. First, a is independent of the income group. The second assumption is that we set $a = 0.5$ which means that if a household expects the change in unemployment to be less than half a percentage point (in absolute terms), she reports that she expects unemployment to be about the same as it is at the time of the survey (our results are quite robust with respect to our choice of a).

As the question in the survey is about the expected change in unemployment, we add the actual unemployment rate at the time of the survey to μ_t^e to construct a time-series of unemployment expectations, as in Bhandari et al. (2019). That said, we will also report the case of expected unemployment *changes*.

Given the so-constructed expectations, we can compute forecast revisions as

$$\mu_t^e - \mu_{t-1}^e$$

and four-quarter-ahead forecast errors using the actual unemployment rate u_t obtained from FRED as

$$u_{t+4} - \mu_t^e. \tag{D.36}$$

For the case of expected unemployment changes, we replace u_{t+4} with $(u_{t+4} - u_t)$ in equation (D.36).

Following Coibion and Gorodnichenko (2015), we then regress forecast errors on forecast revisions

$$u_{t+4} - \mu_t^e = c^e + b^{e,CG}\left(\mu_t^e - \mu_{t-1}^e\right) + \epsilon_t^e, \tag{D.37}$$

to estimate $b^{e,CG}$ for each income group e. Note, however, that the expectations in the forecast revisions are about unemployment at different points in time. To account for this, we instrument forecast revisions by the *main business cycle shock* obtained from Angeletos et al. (2020) (Coibion and Gorodnichenko (2015) use a similar IV strategy when considering expectations from the Michigan Survey).

Table D.1: Regression Results of Equation (D.34)

	IV Regression			OLS		
	Bottom 25%	Middle 50%	Top 25%	Bottom 25%	Middle 50%	Top 25%
$\widehat{b}^{e,CG}$	0.85	0.75	0.63	1.22	1.10	0.90
s.e.	(0.471)	(0.453)	(0.401)	(0.264)	(0.282)	(0.247)
F-stat.	24.76	18.74	17.86	-	-	-
N	152	152	152	157	157	157

Note: This table provides the estimated $\widehat{b}^{e,CG}$ from regression (D.34) for different income groups. The first three columns show the results when the right-hand side in equation (D.34) is instrumented using the *main business cycle shock* from Angeletos et al. (2020) and the last three columns using OLS. Standard errors are robust with respect to heteroskedasticity and are reported in parentheses. The row "F-stat." reports the first-stage F-statistic for the IV regressions.

Table D.1 shows the results. The first three columns report the estimated $b^{e,CG}$ from the IV regressions and the last three columns the same coefficients estimated via OLS. Standard errors are robust with respect to heteroskedasticity and are reported in parentheses. The row "F-stat." reports the first-stage F-statistic for the IV regressions. We see that in all cases $\widehat{b}^{e,CG}$ is positive, suggesting that households of all income groups tend to underreact, consistent with our assumption of $\bar{m} < 1$.

Using equation (D.35) we obtain $\bar{m}^e$ equal to 0.86, 0.87 and 0.88 for the bottom 25%, the middle 50% and the top 25%, respectively for the estimates from the IV regressions and 0.82, 0.83 and 0.85 for the OLS estimates. When estimating $\bar{m}^e$ using expected unemployment *changes* instead of the level, the estimated $\bar{m}^e$ equal 0.57, 0.59 and 0.64 for the IV regressions and 0.77, 0.80 and 0.86 for the OLS regressions.

There are two main take-aways from this empirical exercise: first, it further confirms that $\bar{m} = 0.85$ is a reasonable (but rather conservative) deviation from rational expectations. Second, the data suggests that there is heterogeneity in the degree of rationality conditional on households income. In particular, households with higher income tend to exhibit higher degrees of rationality.[106]

If we consider inflation expectations instead of unemployment expectations, we obtain estimated cognitive discounting parameters of 0.70, 0.75 and 0.78 for the bottom 25%, the middle 50% and the top 25%, respectively. Thus, somewhat lower than for unemployment and the differences across income groups are larger. In particular, higher-income households

[106]This is consistent with other empirical findings on heterogeneous deviations from FIRE. Broer et al. (2022), for example, document that wealthier households tend to have more accurate beliefs, as measured by forecast errors.

tend to be more rational (they discount less) than lower-income households. The differences, however, are overall rather small.

D.3 Figures to Section 4.3

D.3.1 Resolving the Catch-22

We graphically illustrate the Catch-22 (Bilbiie (2021)) of the rational model and the resolution of it in the behavioral HANK model in Figure D.2. Figure D.2 shows the case of *nominal* rate changes. The figure shows on the vertical axis the response of contemporaneous output relative to the initial response in the RANK model with rational expectations for anticipated i.i.d. monetary policy shocks occurring at different times k on the horizontal axis.[107]

The orange-dotted line represents the baseline calibration of the rational HANK model. We see that this model is able to generate contemporaneous amplification of monetary policy shocks, that is, an output response that is relatively stronger than in RANK. Put differently the GE effects amplify the effects of monetary policy shocks. Yet, at the same time, it exacerbates the forward guidance puzzle as shocks occurring in the future have even stronger effects on today's output than contemporaneous shocks.

The black-dashed-dotted line shows how the forward guidance puzzle can be resolved by allowing for $\chi < 1$. Yet, this comes at the cost that the model is unable to generate amplification of contemporaneous monetary policy shocks. Recent empirical findings, however, document that GE effects indeed amplify monetary policy changes (Patterson (2023), Auclert (2019)).

The blue-dashed line shows that the behavioral HANK model, on the other hand, generates both: amplification of contemporaneous monetary policy and a resolution of the forward guidance puzzle, both consistent with the empirical facts.

D.3.2 Stability at the Effective Lower Bound

We illustrate the stability of the behavioral HANK model at the lower bound graphically in Figure D.3. Recall from Section 4.2, the forward-iterated IS equation with a natural rate shock:

$$\widehat{y}_t = -\frac{1}{\gamma}\psi_c \underbrace{\left(\widehat{i}_{ELB} - \widetilde{r}^n\right)}_{>0} \sum_{j=0}^{k} \left(\psi_f + \frac{\kappa}{\gamma}\psi_c\right)^j .$$

Figure D.3 shows the output response in RANK, the rational HANK and the behavioral HANK to different lengths of a binding ELB (depicted on the horizontal axis). The

[107]Under fully-rigid prices (i.e., $\kappa = 0$) the RANK model would deliver a constant response for all k. The same is true for two-agent NK models (TANK), i.e., tractable HANK models without type switching. Whether the constant response would lie above or below its RANK counterpart depends on $\chi \lesseqgtr 1$ in the same way the initial response depends on $\chi \lesseqgtr 1$.

Figure D.2: Resolving the Catch-22

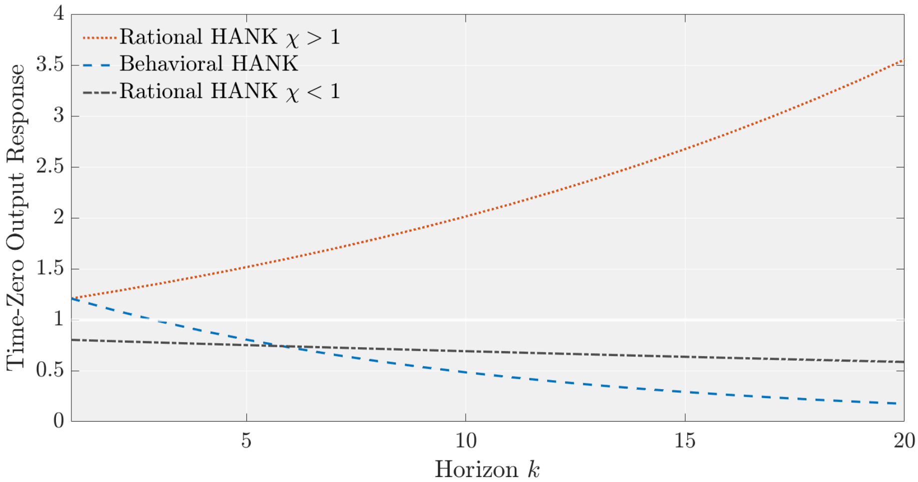

Note: This figure shows the response of total output in period 0 to anticipated i.i.d. monetary policy shocks occurring at different horizons k (horizontal axis), relative to the initial response in the RANK model under rational expectations (equal to 1).

shortcoming of monetary policy due to the ELB, i.e., the gap $\left(\hat{i}_{ELB} - \tilde{r}^n\right) > 0$, is set to a relatively small value of 0.25% (1% annually), and we set $\bar{m} = 0.85$. Figure D.3 shows the implosion of output in the rational RANK (back-solid line) and even more so in the rational HANK model (orange-dotted line): an ELB that is expected to bind for 40 quarters would decrease today's output in the rational RANK by 15% and in the rational HANK model by 45%. On the other hand—and consistent with recent experiences in advanced economies—output in the behavioral HANK model remains quite stable and drops by a mere 3%, as illustrated by the blue-dashed line.

D.4 Extensions and Robustness of the Analytical Model

D.4.1 Robustness of Calibration

In our baseline calibration, we obtain an amplification of conventional monetary policy shocks of 20% compared to the case in which all households are equally exposed to monetary policy (i.e., $\psi_c = 1.2$) for a given share of hand-to-mouth, λ. In particular, we set $\lambda = 0.33$. This results in $\chi = 1.35$. In the quantitative model, we also obtain an amplification of about 20% but this is implied by targeting the micro evidence from Patterson (2023).

To show the robustness of our results, we show in Figure D.4 for different ψ_c (on the horizontal axis) the highest $\bar{m}$ (on the vertical axis) that still resolves the forward-guidance puzzle. The blue-dashed line shows this for $\lambda = 0.33$ and the orange-dotted line for $\lambda = 0.147$ which is the share of borrowing-constrained households in Farhi and Werning (2019).

We see that a $\bar{m}$ of 0.85 (as indicated by the black-solid line) rules out the forward-guidance puzzle in almost all cases. Only at the relatively low λ of 0.147 and a high $\psi_c > 1.48$, we would require a $\bar{m}$ of about 0.84 instead of 0.85 to rule out the forward-guidance puzzle

Figure D.3: The Effective Lower Bound Problem

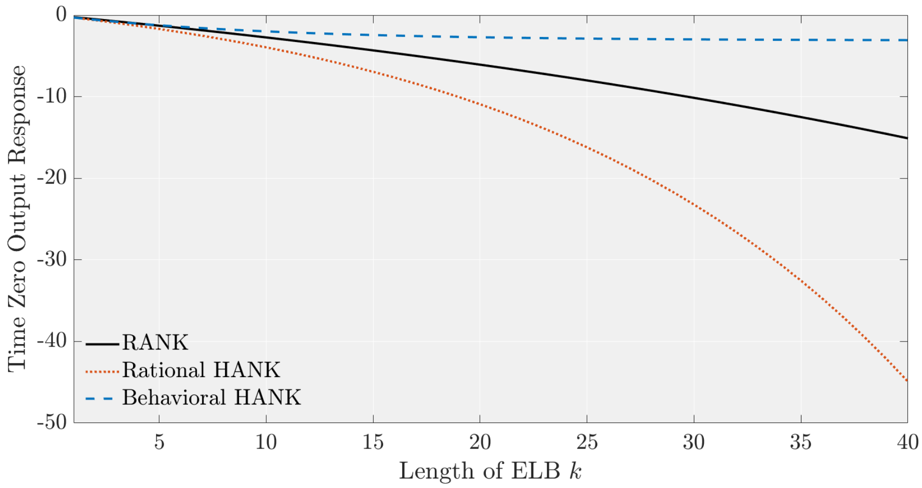

Note: This figure shows the contemporaneous output response for different lengths of a binding ELB k (horizontal axis) and compares the responses across different models.

(note, that a ψ_c of 1.47 at $\lambda = 0.147$ implies $\chi = 2.86$). Given that the empirical estimates point towards values of $\bar{m} \in [0.6, 0.85]$, we conclude the resolution of the forward guidance puzzle in the behavioral HANK model with countercyclical income risk is quite robust.

Figure D.4: Robustness of Forward Guidance Puzzle Solution

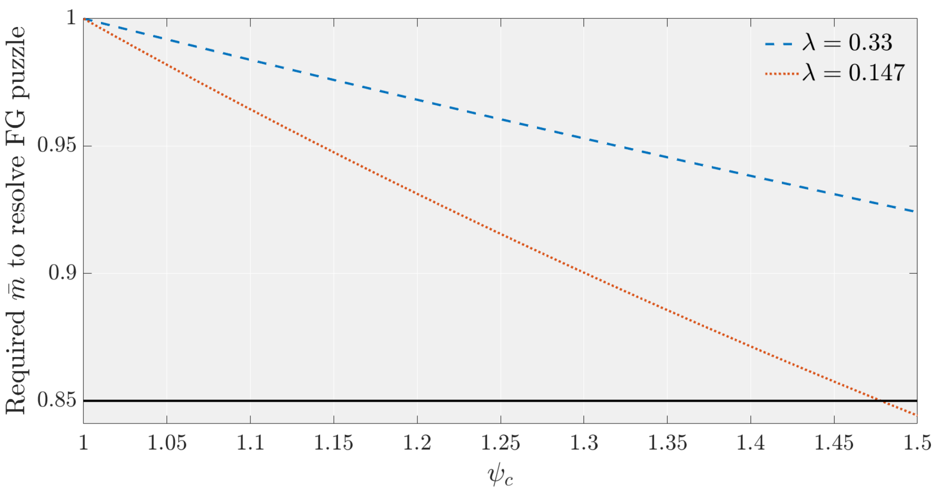

Note: This figures show for different ψ_c (horizontal axis) the required $\bar{m}$ to resolve the forward-guidance puzzle on the vertical axis. The blue-dashed line shows this for our benchmark calibration of $\lambda = 0.33$ and the orange-dotted line for $\lambda = 0.147$.

Also note that the values for γ and κ that we use are directly taken from Bilbiie (2021,0) and are quite standard in the literature. Gabaix (2020), however, sets $\kappa = 0.11$ and $\gamma = 5$. Even though these coefficients differ quite substantially from our baseline calibration, note that our results would barely be affected by this. To see this, note that *amplification* is only determined by λ and χ, both independent of κ and γ. The determinacy condition on the other hand depends on both, κ and γ, but what ultimately matters is the fraction $\frac{\kappa}{\gamma}$ (see Proposition 7). As κ and γ are both approximately five times larger in Gabaix (2020) compared to Bilbiie (2021) and our baseline calibration, the fraction is approximately the same and thus, the determinacy region under an interest-rate peg remains unchanged.

D.4.2 Nominal Interest Rate Changes

In Section 4.3, we focused on the case where monetary policy directly controls the *real* rather than the nominal interest rate. We now show that our results are unchanged when instead focusing on nominal rate changes. As in the main text, we consider two different monetary policy experiments: (i) a contemporaneous monetary policy shock, i.e., a surprise decrease in the nominal interest rate today, and (ii) a forward guidance shock, i.e., a news shock today about a decrease in the nominal interest rate k periods in the future. In both cases, we focus on *i.i.d.* shocks and the Taylor response coefficient is zero, $\phi = 0$.[108]

Proposition 8 *In the behavioral HANK model, there is amplification of contemporaneous monetary policy relative to RANK if and only if*

$$\psi_c > 1 \Leftrightarrow \chi > 1, \tag{D.38}$$

and the forward guidance puzzle is ruled out if

$$\psi_f + \frac{\kappa}{\gamma}\psi_c < 1. \tag{D.39}$$

We thus see, that the amplification result is unchanged (see Proposition 6) whereas the condition to rule out the forward-guidance puzzle is somewhat stricter as $\frac{\kappa}{\gamma}\psi_c > 0$. This is the case because there is now an inflation feedback effect. An expected decrease in the nominal interest rate in the future increases inflation expectations and thus, lowers the real rate further. Thus, the effects on today's output become stronger.

However, again a relatively small underreaction of the behavioral households is enough to resolve the forward guidance puzzle. Given our calibration there is no forward guidance puzzle in the behavioral HANK model as long as $\bar{m} < 0.94$ which is above the upper bounds for empirical estimates (see Section 4.2).

D.4.3 Allowing for Steady State Inequality

In the tractable model, we have assumed that there is no steady state inequality, i.e., $C^H = C^U$. In the following, we relax this assumption and denote steady state inequality by $\Omega \equiv \frac{C^U}{C^H}$. Recall the Euler equation of unconstrained households

$$\left(C_t^U\right)^{-\gamma} = \beta R_t \mathbb{E}_t^{BR}\left[s\left(C_t^U\right)^{-\gamma} + (1-s)\left(C_t^H\right)^{-\gamma}\right],$$

[108] If we instead impose $\phi > 0$, contemporaneous amplification in the following proposition is not affected but the condition to rule out the forward guidance puzzle is further relaxed.

from which we can derive the steady state real rate

$$R = \frac{1}{\beta(s + (1-s)\Omega^{\gamma})}.$$

Log-linearizing the Euler equation yields

$$\widehat{c}_t^U = \beta R \bar{m} \left[s\mathbb{E}_t \widehat{c}_{t+1}^U + (1-s)\Omega^{\gamma}\mathbb{E}_t \widehat{c}_{t+1}^H \right] - \frac{1}{\gamma}\left(\widehat{i}_t - \mathbb{E}_t \pi_{t+1} \right).$$

Combining this with the consumption functions and the steady state real rate yields the IS equation

$$\widehat{y}_t = \bar{m}\tilde{\delta}\mathbb{E}_t \widehat{y}_{t+1} - \frac{1}{\gamma}\frac{1-\lambda}{1-\lambda\chi}\left(\widehat{i}_t - \mathbb{E}_t \pi_{t+1} \right), \tag{D.40}$$

with

$$\tilde{\delta} \equiv 1 + (\chi - 1)\frac{(1-s)\Omega^{\gamma}}{s + (1-s)\Omega^{\gamma}}\frac{1}{1-\lambda\chi}.$$

From a qualitative perspective, the whole analysis in Section 4.3 could be carried out with $\tilde{\delta}$ instead of δ. Quantitatively the differences are small as well. For example, if we set $\Omega = 1.5$, we get $\tilde{\delta} = 1.05$ instead of $\delta = 1.034$. Thus, we need $\bar{m} < 0.93$ instead of $\bar{m} < 0.94$ for determinacy under a peg.

D.4.4 Persistent Monetary Policy Shocks

In the main text in Section 4.3, we illustrated the resolution of the Catch-22 by considering i.i.d. monetary policy shocks (following Bilbiie (2021)). The behavioral HANK model delivers initial amplification of these monetary shocks but the effects decrease with the horizon of the shock, i.e., the behavioral HANK model resolves the forward guidance puzzle. Another way to see this is by considering persistent shocks.

Figure D.5 illustrates this. The figure shows the response of output in period t to a shock in period t for different degrees of persistence (x-axis). The black-solid line shows the output response in RANK and the blue-dashed line in the behavioral HANK. The forward guidance puzzle in RANK manifests itself in the sense that highly persistent shocks have stronger effects in RANK than in the behavioral HANK. Persistent shocks are basically a form of forward guidance and thus, with high enough persistence in the shocks, the RANK model predicts stronger effects than the behavioral HANK model.

As the persistence of the monetary policy shock approaches unity, the rational model leads to the paradoxical finding that an exogenous increase in the nominal interest rate leads to an expansion in output. To see this, note that we can write output as

$$\widehat{y}_t = -\frac{\frac{\psi_c}{\gamma}}{1 + \frac{\psi_c}{\gamma}\phi\kappa - \left(\psi_f + \psi_c \frac{\kappa}{\gamma} \right)\rho}\varepsilon_t^{MP}. \tag{D.41}$$

Given our baseline calibration and a Taylor coefficient of $\phi = 1$, the rational model would

Figure D.5: Initial Output Response for Varying Degrees of the Persistence

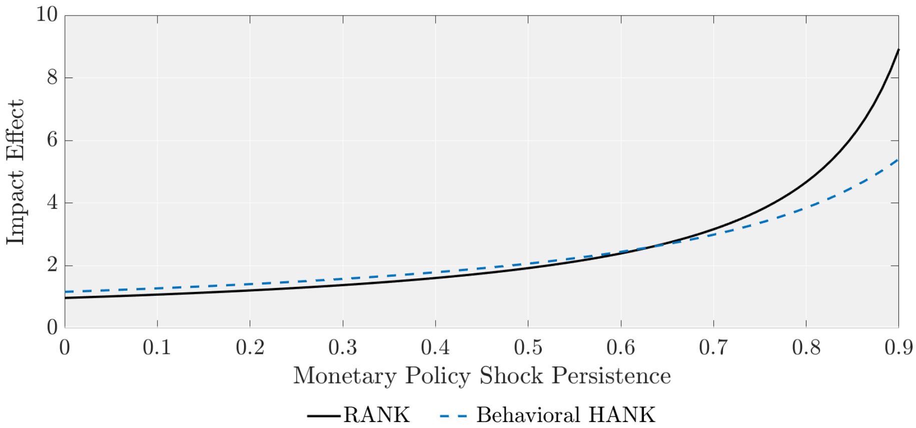

Note: This figure shows the initial output response to monetary policy shocks with different degrees of persistence.

produce these paradoxical findings for $\rho > 0.967$. The behavioral HANK model, on the other hand, does not suffer from this as the denominator is always positive, even when $\phi = 0$ and $\rho = 1$.

D.4.5 Forward-Looking NKPC and Real Interest Rates

In the tractable model, we made the assumption that agents are rational with respect to real interest rates (as in Gabaix (2020)) and assumed a static Phillips Curve for simplicity. We now show that the results are barely affected when considering a forward-looking New Keynesian Phillips Curve (NKPC) and that agents are also boundedly rational with respect to real rates. Gabaix (2020) derives the NKPC under bounded rationality and shows that it takes the form:

$$\pi_t = \beta M^f \mathbb{E}_t \pi_{t+1} + \kappa \widehat{y}_t,$$

with

$$M^f \equiv \bar{m} \left(\theta + \frac{1 - \beta\theta}{1 - \beta\theta\bar{m}} (1 - \theta) \right),$$

where $1 - \theta$ captures the Calvo probability of price adjustment.

Taking everything together (including the bounded rationality with respect to real interest rates), the model can be summarized by the following three equations:

$$\begin{aligned}
\widehat{y}_t &= \psi_f \mathbb{E}_t \widehat{y}_{t+1} - \psi_c \frac{1}{\gamma} \left(\widehat{i}_t - \bar{m} \mathbb{E}_t \pi_{t+1} \right) \\
\pi_t &= \beta M^f \mathbb{E}_t \pi_{t+1} + \kappa \widehat{y}_t \\
\widehat{i}_t &= \phi \pi_t.
\end{aligned}$$

Plugging the Taylor rule into the IS equation, we can write everything in matrix form:

$$\begin{pmatrix} \mathbb{E}_t\pi_{t+1} \\ \mathbb{E}_t\widehat{y}_{t+1} \end{pmatrix} = \underbrace{\begin{pmatrix} \frac{1}{\beta M^f} & -\frac{\kappa}{\beta M^f} \\ \frac{\psi_c}{\gamma\psi_f}\left(\phi - \frac{\bar{m}}{\beta M^f}\right) & \frac{1}{\psi_f}\left(1 + \frac{\psi_c \bar{m}\kappa}{\gamma\beta M^f}\right) \end{pmatrix}}_{\equiv A} \begin{pmatrix} \pi_t \\ \widehat{y}_t \end{pmatrix}. \tag{D.42}$$

For determinacy, we need

$$det(A) > 1; \quad det(A) - tr(A) > -1; \quad det(A) + tr(A) > -1.$$

The last condition is always satisfied. The first two conditions are satisfied if and only if

$$\phi > \max\left\{ \frac{\beta\delta M^f \bar{m} - 1}{\frac{\kappa}{\gamma}\frac{1-\lambda}{1-\lambda\chi}}, \bar{m} + \frac{(\delta\bar{m} - 1)(1 - \beta M^f)}{\frac{\kappa}{\gamma}\frac{1-\lambda}{1-\lambda\chi}} \right\}.$$

In the case of a static Phillips curve but bounded rationality with respect to the real rate, the second condition is the crucial one. To capture the static Phillips curve, we can simply set $M^f = 0$. We can see that bounded rationality with respect to the real rate relaxes the determinacy condition whereas a forward-looking NKPC tightens it. But even in the case of a forward-looking NKPC (rational or behavioral), cognitive discounting relaxes the determinacy condition and thus, all our results from the static Phillips curve are qualitatively unchanged. Under our baseline calibration and $\theta = 0.875$ and $\beta = 0.99$ as in Gabaix (2020), the model still features determinacy under a peg, even when real interest rate expectations are rational (and therefore, also when they are behavioral).

D.4.6 Cognitive Discounting of the State Vector

In Section 4.2, we assume that cognitive discounting applies to all variables, which differs slightly from the assumption in Gabaix (2020) who assumes that cognitive discounting applies to the *state* of the economy (exogenous shocks as well as announced monetary and fiscal policies). He then proves (Lemma 1 in Gabaix (2020)) how cognitive discounting applies as a result (instead of as an assumption) to all future variables, including future consumption choices. For completeness, we show in this section how our results are unaffected when following the approach in Gabaix (2020).

Let X_t denote the (de-meaned) state vector which evolves as

$$X_{t+1} = G^X\left(X_t, \varepsilon_{t+1}\right), \tag{D.43}$$

where G^X denotes the transition function of X in equilibrium and ε are zero-mean innovations. Linearizing equation (D.43) yields

$$X_{t+1} = \Gamma X_t + \varepsilon_{t+1}, \tag{D.44}$$

where ε_{t+1} might have been renormalized. The assumption in Gabaix (2020) is that the behavioral agent perceives the state vector to follow

$$X_{t+1} = \bar{m} G^X(X_t, \varepsilon_{t+1}), \tag{D.45}$$

or in linearized terms

$$X_{t+1} = \bar{m}\left(\Gamma X_t + \varepsilon_{t+1}\right). \tag{D.46}$$

The expectation of the boundedly-rational agent of X_{t+1} is thus $\mathbb{E}_t^{BR}\left[X_{t+1}\right] = \bar{m}\mathbb{E}_t\left[X_{t+1}\right] = \bar{m}\Gamma X_t$. Iterating forward, it follows that $\mathbb{E}_t^{BR}\left[X_{t+k}\right] = \bar{m}^k \mathbb{E}_t\left[X_{t+k}\right] = \bar{m}^k \Gamma^k X_t$.

Now, consider any variable $z(X_t)$ with $z(0) = 0$ (e.g., demeaned consumption of unconstrained households $C^U(X_t)$). Linearizing $z(X)$, we obtain $z(X) = b_X^z X$ for some b_X^z and thus

$$\begin{aligned}
\mathbb{E}_t^{BR}\left[z(X_{t+k})\right] &= \mathbb{E}_t^{BR}\left[b_X^z X_{t+k}\right] = b_X^z \mathbb{E}_t^{BR}\left[X_{t+k}\right] \\
&= b_X^z \bar{m}^k \mathbb{E}_t\left[X_{t+k}\right] = \bar{m}^k \mathbb{E}_t\left[b_X^z X_{t+k}\right] \\
&= \bar{m}^k \mathbb{E}_t\left[z(X_{t+k})\right].
\end{aligned}$$

For example, expected consumption of unconstrained households tomorrow (in linearized terms) is given by

$$\mathbb{E}_t^{BR}\left[\widehat{c}^U(X_{t+1})\right] = \bar{m}\mathbb{E}_t\left[\widehat{c}^U(X_{t+1})\right], \tag{D.47}$$

which we denote in the main text as

$$\mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^U\right] = \bar{m}\mathbb{E}_t\left[\widehat{c}_{t+1}^U\right]. \tag{D.48}$$

Now, take the linearized Euler equation (4.16) of unconstrained households:

$$\widehat{c}_t^U = s\mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^U\right] + (1-s)\mathbb{E}_t^{BR}\left[\widehat{c}_{t+1}^H\right] - \frac{1}{\gamma}\widehat{r}_t, \tag{D.49}$$

where $\widehat{r}_t \equiv \widehat{i}_t - \mathbb{E}_t \pi_{t+1}$.

Using the notation in Gabaix (2020), we can write the Euler equation as

$$\widehat{c}^U(X_t) = s\mathbb{E}_t^{BR}\left[\widehat{c}^U(X_{t+1})\right] + (1-s)\mathbb{E}_t^{BR}\left[\widehat{c}^H(X_{t+1})\right] - \frac{1}{\gamma}\widehat{r}(X_t). \tag{D.50}$$

Now, applying the results above, we obtain

$$\widehat{c}^U(X_t) = s\bar{m}\mathbb{E}_t\left[\widehat{c}^U(X_{t+1})\right] + (1-s)\bar{m}\mathbb{E}_t\left[\widehat{c}^H(X_{t+1})\right] - \frac{1}{\gamma}\widehat{r}(X_t), \tag{D.51}$$

which after writing $\widehat{c}^U(X_t)$, $\widehat{c}^U(X_{t+1})$ and $\widehat{c}^H(X_{t+1})$ in terms of total output yields exactly the IS equation in Proposition 5.

D.4.7 Microfounding $\bar{m}$

Gabaix (2020) shows how to microfound $\bar{m}$ from a noisy signal extraction problem in the case of a representative agent. Following these lines, we show how such a signal-extraction problem offers a potential microfoundation in the heterogeneous agent case, too.

The (linearized) law of motion of the state variable, X_t, is given by $X_{t+1} = \Gamma X_t + \varepsilon_{t+1}$ (a similar reasoning extends to the non-linearized case), where X has been demeaned. Now assume that each households j performs a mental simulation of the future, but receives only noisy signals about that simulation, i.e., the household receives signals S^j_{t+1} of X_{t+1}, and these signals are given by

$$S^j_{t+1} = \begin{cases} X_{t+1} & \text{with probability } p \\ X'_{t+1} & \text{with probability } 1-p \end{cases}$$

where X'_{t+1} is an i.i.d. draw from the unconditional distribution of X_{t+1}, which has an unconditional mean of zero. In words, with probability p the agent j receives perfectly precise information in one particular mental simulation of the future, and with probability $1-p$ agent j receives a signal realization that is completely uninformative. A fully-informed rational agent would have $p = 1$.

The household runs a continuum of these simulations in his head. The conditional mean of X_{t+1}, given the signal S^j_{t+1}, is given by

$$X^e_{t+1} \equiv \mathbb{E}\left[X_{t+1} | S_{t+1} = s^j_{t+1}\right] = p \cdot s^j_{t+1}.$$

To see this, note that the joint distribution of (X_{t+1}, S^j_{t+1}) is

$$f(x_{t+1}, s^j_{t+1}) = pg(s^j_{t+1})\delta_{s^j_{t+1}}(x_{t+1}) + (1-p)g(s^j_{t+1})g(x_{t+1}),$$

where $g(X_{t+1})$ denotes the distribution of X_{t+1} and δ is the Dirac function. Given that the unconditional mean of X_{t+1} is 0, i.e., $\int x_{t+1} g(x_{t+1}) dx_{t+1} = 0$, it follows that

$$\begin{aligned} \mathbb{E}_t\left[X_{t+1} | S^j_{t+1} = s^j_{t+1}\right] &= \frac{\int x_{t+1} f(x_{t+1}, s^j_{t+1}) dx_{t+1}}{\int f(x_{t+1}, s^j_{t+1}) dx_{t+1}} \\ &= \frac{pg(s^j_{t+1})s^j_{t+1} + (1-p)g(s^j_{t+1}) \int x_{t+1} g(x_{t+1}) dx_{t+1}}{g(s^j_{t+1})} \\ &= ps^j_{t+1}. \end{aligned}$$

Furthermore, we have

$$\mathbb{E}\left[S_{t+1} | X_{t+1}\right] = pX_{t+1} + (1-p)\mathbb{E}\left[X'_{t+1}\right] = pX_{t+1}.$$

So, it follows that the *average* expectation of X_{t+1} over all these simulations is given by

$$\begin{aligned}\mathbb{E}\left[X^e_{t+1}(S_{t+1})|X_{t+1}\right] &= \mathbb{E}\left[p \cdot S_{t+1}|X_{t+1}\right] \\ &= p \cdot \mathbb{E}\left[S_{t+1}|X_{t+1}\right] \\ &= p^2 X_{t+1}.\end{aligned}$$

Defining $\bar{m} \equiv p^2$ and since $X_{t+1} = \Gamma X_t + \varepsilon_{t+1}$, we have that the agent perceives the law of motion of X to equal

$$X_{t+1} = \bar{m}\left(\Gamma X_t + \varepsilon_{t+1}\right), \tag{D.52}$$

as imposed in equation (D.46). The boundedly-rational expectation of X_{t+1} is then given by

$$\mathbb{E}^{BR}_t\left[X_{t+1}\right] = \bar{m}\mathbb{E}_t\left[X_{t+1}\right].$$

D.5 Details and Extensions to Section 4.4

D.5.1 Robustness of Calibration

Following Patterson (2023), we calibrate the unequal income exposure of households such that a linear regression of the income elasticity w.r.t. GDP on MPCs yields a coefficient of 1.33. To show that our results are robust to more extreme calibrations, Figure D.6 shows the case where we target a coefficient of 2.0. As one would expect, contemporaneous monetary policy shocks are further amplified and due to the induced countercyclical income risk forward guidance becomes somewhat more effective. Overall, however, we conclude that the forward-guidance puzzle is still clearly ruled out and our results thus robust.

Figure D.6: Robustness of Calibration

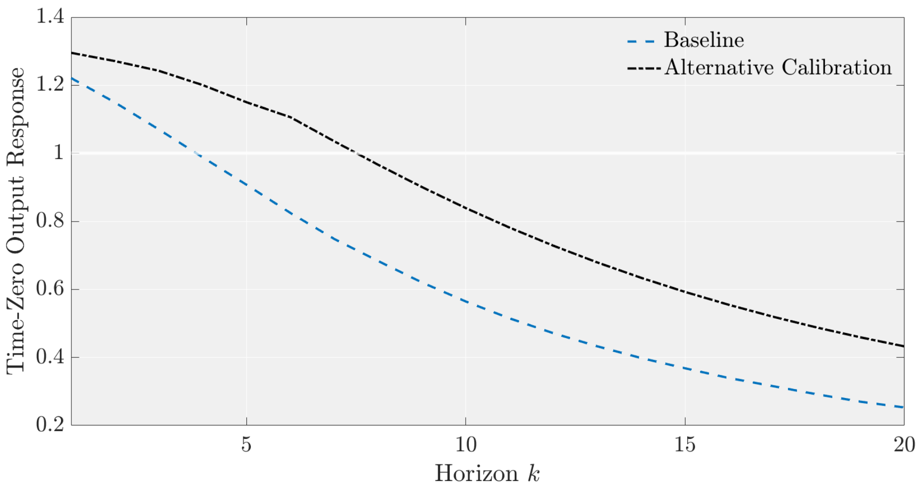

Note: This figure shows the response of total output in period 0 to anticipated i.i.d. monetary policy shocks occurring at different horizons k for a more unequal income exposure of households.

D.5.2 Stability at the ELB and Fiscal Multipliers

Figure D.7 shows the output and nominal interest rate response after a shock to the discount factor in the quantitative behavioral HANK model and in its rational counterpart. In particular, the discount factor jumps on impact by 0.65% for 12 quarters before it returns to steady state.

Figure D.7: ELB recession in the quantitative behavioral HANK model

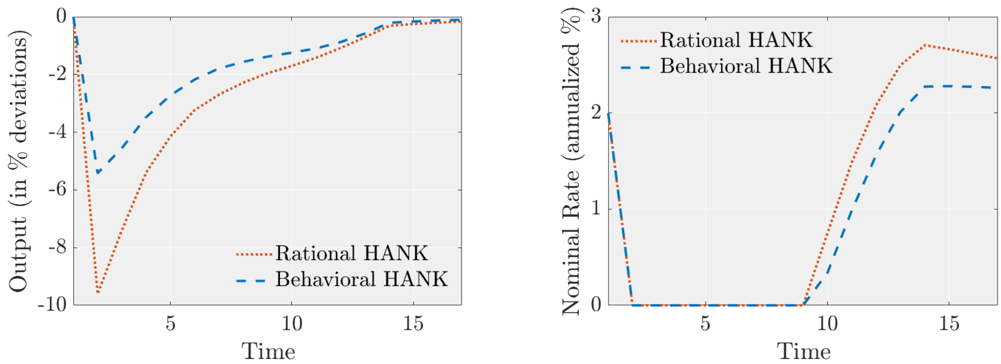

Note: This figure shows the impulse responses of total output and of the nominal interest rate after a discount factor shock that brings the economy to the ELB for 8 quarters.

We see that while the interest-rate path is quite similar across the two models, the output drop in the rational model is about twice as deep as in the behavioral HANK model. The intuition is as in the tractable model (Section 4.3). The binding ELB acts like a contractionary monetary policy shock because the nominal interest rate cannot keep up with the drop in the natural rate due to the ELB. Under rational expectations, households fully account for this and thus, cut back their consumption quite strongly on impact. Thus, the ELB leads to a large recession. Under cognitive discounting, on the other hand, households discount these future shocks and hence, decrease their consumption by less, leading to a milder recession.

D.5.3 Heterogeneous $\bar{m}$: Alternative Calibration

The estimated differences in households' underreaction across different income groups are rather small. Nevertheless, one might argue that some agents (financial markets, for example) closely track what the Fed is doing and that they are usually well informed about its actions. To mirror this, we assume that the highest-productivity households are fully rational, i.e., their $\bar{m}$ is equal to 1. To keep the average $\bar{m}$ at 0.85, we then assume that the lowest-productivity households have a $\bar{m}$ of 0.7 and the middle-productivity households of 0.85.

The black-dashed-dotted line in Figure D.8 shows the time zero output response (vertical axis) to an announced monetary policy shock taking place at different horizons (horizontal axis).

We see that forward guidance is more powerful than in the baseline calibration as the agents that tend to be more forward looking because they are not at their borrowing constraint are also more rational. Overall, however, our results remain robust. Thus, even when a subpopulation of all households is fully rational, the behavioral HANK model can simultaneously generate amplification of conventional monetary policy through indirect effects and rule out the forward guidance puzzle.

Figure D.8: Heterogeneous $\bar{m}$ and Monetary Policy

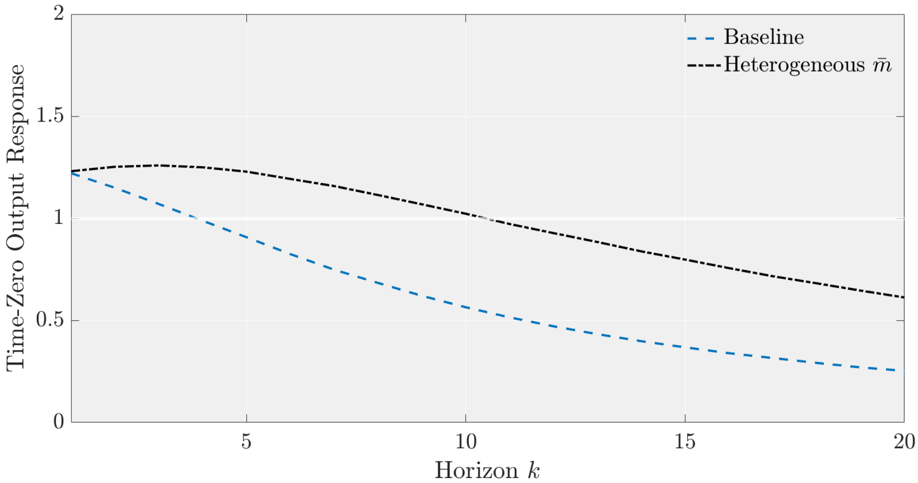

Note: This figure shows the response of total output in period 0 to anticipated i.i.d. monetary policy shocks occurring at different horizons k for the baseline calibration with $\bar{m} = 0.85$ for all households (blue-dashed line), and for the model in which high productivity households have $\bar{m} = 1$, medium-level productivity households have $\bar{m} = 0.85$ and low-productivity households have $\bar{m} = 0.7$ (black-dashed-dotted line).

D.6 Additional Results and Figures to Section 4.5

D.6.1 Decomposition of Amplification Channel: More Cognitive Discounting

As shown in Section 4.5, household heterogeneity and cognitive discounting interact in such a way that productivity shocks get amplified through both of these ingredients as well as their interaction. Given our baseline calibration, the interaction accounts for about 19% of the additional increase compared to RANK. We now consider an alternative calibration where we set the cognitive discounting parameter $\bar{m}$ to 0.6 instead of 0.85. Thus, somewhat closer to the lower bound of empirical estimates (see Section 4.2). Figure D.9 shows the decomposition of the additional amplification of negative productivity shocks under this alternative calibration.

Two things stand out. First, the overall inflation increase is more than twice as large compared to RANK. Given our discussion in Section 4.5, this is no surprise. The stronger cognitive discounting induces a larger increase in inflation after the negative productivity shock. Second, the interaction becomes even more important. In fact, the interaction alone

Figure D.9: Decomposition of the Additional Inflation Increase: Lower $\bar{m}$

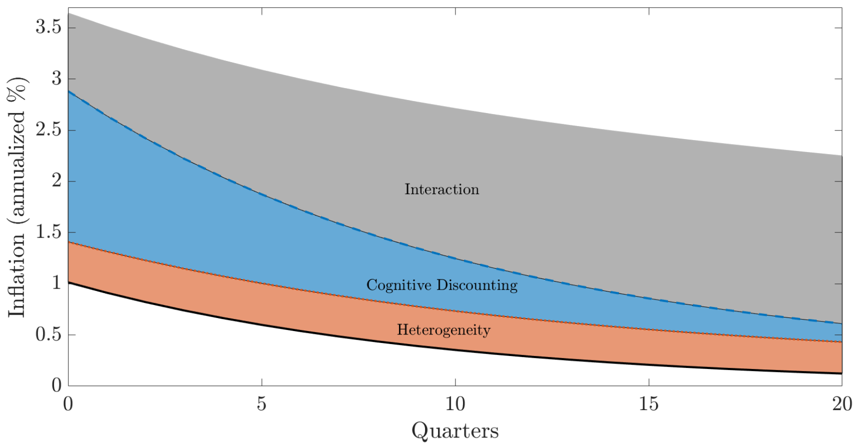

Note: This figure shows the decomposition of the additional inflation increase in the behavioral HANK model for $\bar{m} = 0.6$ compared to the rational RANK model. The orange-shaded area represents the additional increase that is solely due to the heterogeneous exposure of households, the blue area the increase due to cognitive discounting and the gray area the additional increase that is due to the interaction of heterogeneity and cognitive discounting.

accounts for more than the underlying heterogeneity itself. It amounts to more than 75% of the impact inflation response in RANK (1 percentage point) or about 29% of the *additional increase.*

D.6.2 Procyclical HANK

Figure D.10: Procyclical inequality

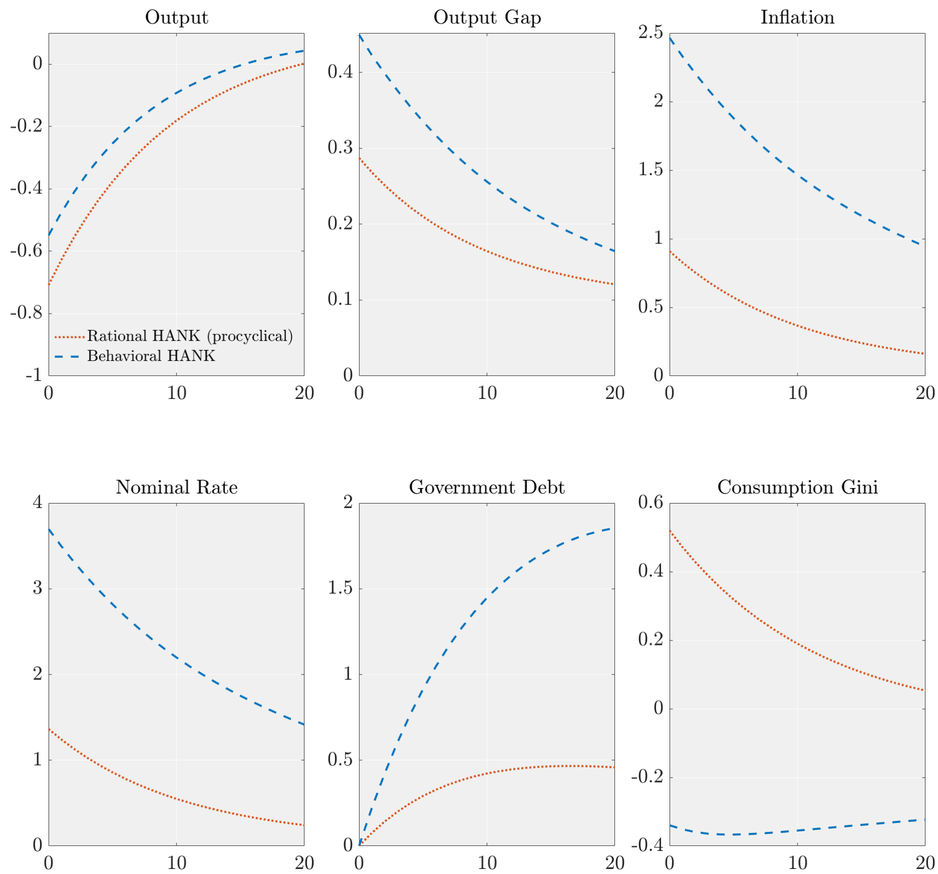

Note: This figure shows the impulse responses after a TFP shock that decreases potential output by 1% when monetary policy follows a Taylor rule for the behavioral HANK model (blue-dashed lines) and for the rational HANK model with procyclical inequality (orange-dotted lines). Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

D.6.3 Behavioral Firms

Figure D.11 shows the impulse-response functions after a negative productivity shock when monetary policy follows a Taylor rule and in which firms are behavioral (with a cognitive discounting factor of 0.85). We see that the increase in inflation when monetary policy follows a Taylor rule is somewhat muted whereas the increase in the output gap is strongly amplified compared to the case in which firms are rational. The reason is that firms discount the increase in their future marginal costs and thus increase their prices not as strongly. According to the Taylor rule this then leads to a smaller increase in the nominal interest rate (both channels inducing a lower real rate) so that households consume more, leading to an increase in demand and thus, the output gap.

Figure D.11: Inflationary supply shock: behavioral firms

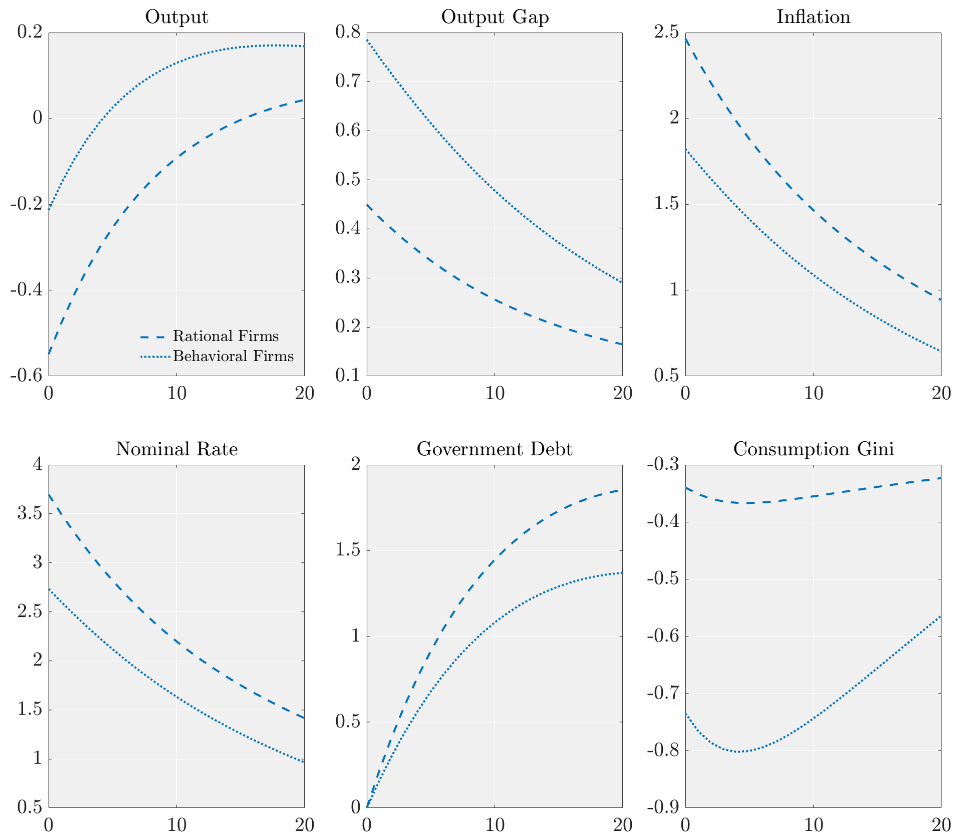

Note: This figure shows the impulse responses after a productivity shock for the case that monetary policy follows a Taylor rule and firms cognitively discount the future with a cognitive discounting parameter of 0.85. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per annual-GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

D.6.4 Cost-Push Shocks

We now show that the fiscal and monetary implications are very similar for an inflationary cost-push shock. To introduce cost-push shocks, we assume that the desired mark-up of firms, μ_t follows an AR(1)-process, $\mu_t = (1-\rho_\mu)\bar{\mu} + \rho_\mu \mu_{t-1} + \varepsilon_t^\mu$, where ε_t^μ is an i.i.d. shock, $\bar{\mu}$ the steady-state level of the desired markup and ρ_μ the persistence of the shock process which we set to $\rho_\mu = 0.9$. The rest of the model is as in Section 4.5. Note, that we model the shock such that it also applies to the model under flexible prices, thus moves potential output as well.

Figure D.12 shows the impulse-response functions of output, the output gap, inflation, nominal interest rates, government debt and the consumption Gini index as a measure of consumption inequality following an inflationary cost-push shock. The blue-dashed lines show the responses in the behavioral HANK model with homogeneous and heterogeneous $\bar{m}$, respectively, and the orange-dotted lines in the rational HANK model. In both cases, monetary policy fully stabilizes inflation by assumption. Output drops, with the responses being practically identical across the two models. Again, the output gap is practically closed in both models. The required response of the nominal interest rate, however, differs substantially across the behavioral and the rational model, as was the case after a negative productivity shock, discussed in Section 4.5. In the behavioral HANK model the monetary authority increases the nominal rate much more strongly and more persistently. The reason for this strong response is that households cognitively discount future (expected) interest rate hikes making them less effective for stabilizing inflation today. Thus, in order to achieve the same stabilization outcome in every period, the interest rate needs to increase by more.

Increasing the interest rate more strongly increases the cost of debt for the government which it finances in the short run by issuing more debt. The middle panel on the bottom line in Figure D.12 shows that government debt in the behavioral model increases more than three times as much as in the rational model. Furthermore, consumption inequality increases in both models. There are two channels: first and most important, the cost-push shock increases dividends and decreases wages which redistributes from low to high productivity households thereby pushing up consumption inequality. Second, the increase in the real interest rate redistributes towards high wealth households but it is the high productivity households who eventually pay the tax burden. This slightly decreases the consumption of high productivity households and increases the consumption of middle productivity households who hold some assets but do not face tax increases. Thus, the second channel slightly dampens the increase in inequality and, as real interest rates increase by more, this channel is stronger in the behavioral HANK model.

Figure D.13 shows the impulse-response functions of output, the output gap, inflation, nominal interest rates, government debt (as a share of annual GDP) and consumption inequality for the same cost-push shock but for the case in which monetary policy follows a simple Taylor rule with a response coefficient of 1.5.

Figure D.12: Inflationary cost-push shock: strict inflation targeting

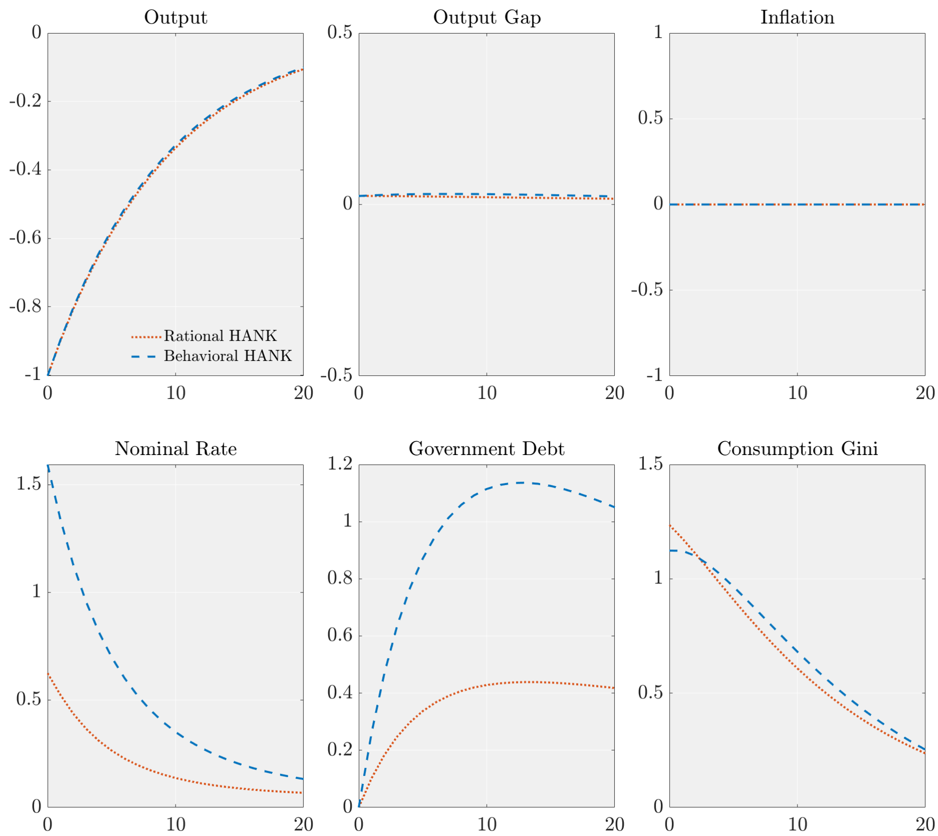

Note: This figure shows the impulse responses after a cost-push shock that decreases potential output by 1% in the inflation-stabilizing monetary policy regime. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

As in the case where monetary policy fully stabilizes inflation, inflation and the nominal interest rate increase substantially more strongly in the behavioral HANK model than in its rational version. Also government debt increases more substantially.

Consumption inequality increases less strongly than with fully stabilizing inflation. The overheating economy—reflected in the positive output gap and increase in inflation—increases wages and decreases profits (relative to the inflation stabilizing regime) in the same way as expansionary policy shocks in Sections 4.3 and 4.4 do, thereby redistributing towards lower income households which dampens the increase in consumption inequality.

Figure D.13: Inflationary cost-push shock: Taylor rule

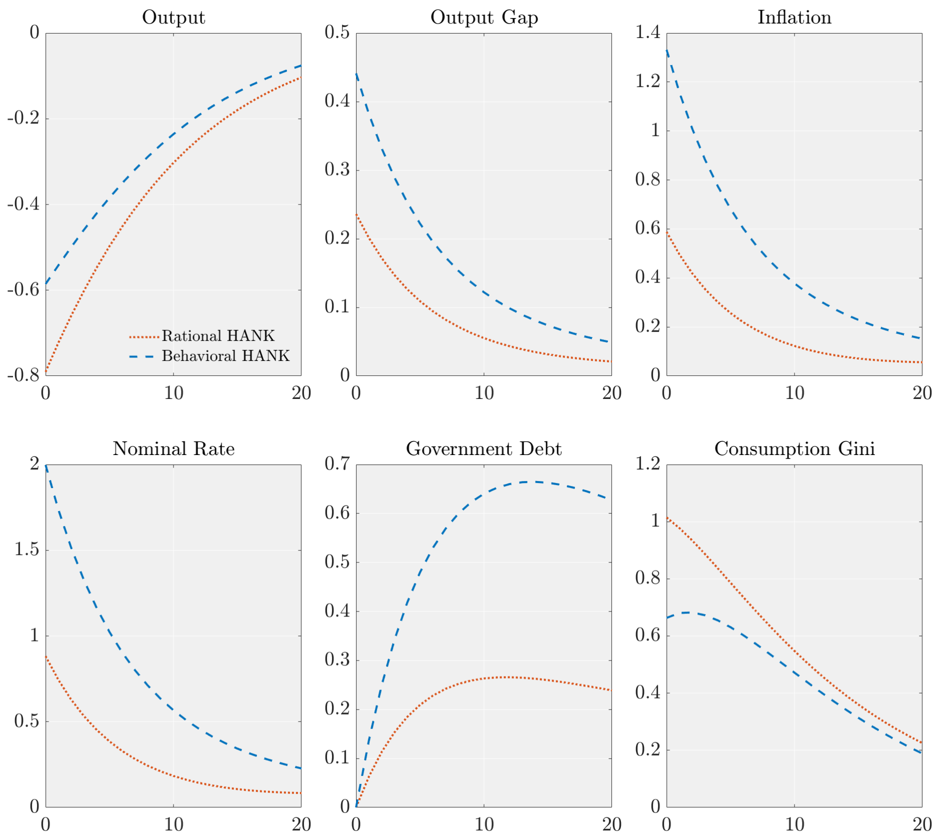

Note: This figure shows the impulse responses after a cost-push shock that decreases potential output by 1% in the Taylor rule monetary policy regime. Output and the output gap are shown as percentage deviations from steady state output, the nominal interest rate and inflation as annualized percentage points and the government debt level as percentage point deviations of the debt-per-annual GDP level. The lower-right figure shows the change in the consumption Gini index as a percentage deviation from the stationary equilibrium.

Appendix E

Appendix for Chapter 5

E.1 Additional Results

Table E.1: Subjective financial condition forecasts are strongly positively correlated with income forecasts

	Forecasted probability of increase in:			
	Nominal income		Real income	
	Unweighted	Weighted	Unweighted	Weighted
	(1)	(2)	(3)	(4)
1= Optimistic forecast of sfc	0.00487	0.00484	0.00576	0.00546
s.e.	(0.00015)	(0.00020)	(0.00018)	(0.00024)
N	15,047	15,047	15,049	15,049
N panelists	3057	3057	3056	3056

Notes: Each column presents results from a single OLS regression of the row variable on the column variable and a constant. Standard errors, clustered on panelist, in parentheses. Weighted estimates use the ALP sampling probability weight for each observation. Income forecasts in percentage point units, so e.g., a point estimate of 0.005 indicates a 1/2 percentage point increase in sfc optimism per 1 pp increase in the probability of an income increase. SFC forecast optimism is indicated by responding to the question "Now looking ahead - do you think that a year from now you will be better off financially, or worse off, or just about the same as now?" with "Will be better off".

Table E.2: Household financial condition forecasts and forecast errors tilt optimistic

Panel A. All forecasts, unweighted	Realization this year			
Forecast last year	Better	Same	Worse	Total
Better	0.10	0.13	0.04	0.27
Same	0.06	0.45	0.10	0.61
Worse	0.01	0.05	0.07	0.12
Total	0.16	0.63	0.21	1
Panel B. July 2009 & 2010, unweighted	Realization this year			
Forecast last year	Better	Same	Worse	Total
Better	0.06	0.16	0.05	0.28
Same	0.05	0.40	0.15	0.60
Worse	0.01	0.05	0.07	0.12
Total	0.12	0.61	0.27	1
Panel C. July 2009 & 2010, weighted	Realization this year			
Forecast last year	Better	Same	Worse	Total
Better	0.07	0.18	0.05	0.30
Same	0.04	0.38	0.14	0.56
Worse	0.01	0.07	0.06	0.14
Total	0.12	0.63	0.25	1
Panel D. January 2015 & 2016, unweighted	Realization this year			
Forecast last year	Better	Same	Worse	Total
Better	0.10	0.14	0.04	0.28
Same	0.06	0.47	0.08	0.61
Worse	0.01	0.05	0.06	0.12
Total	0.17	0.66	0.18	1
Panel E. January 2015 & 2016, weighted	Realization this year			
Forecast last year	Better	Same	Worse	Total
Better	0.11	0.13	0.03	0.27
Same	0.05	0.50	0.08	0.63
Worse	0.01	0.04	0.05	0.10
Total	0.17	0.67	0.16	1

Note: Cells report sample proportions. Forecasts: "Now looking ahead - do you think that a year from now you will be better off financially, or worse off, or just about the same as now?" Response options: Will be better off/About the same/Will be worse off. Realizations: "We are interested in how people are getting along financially these days. Would you say that you are better off or worse off financially than you were a year ago?" Response options: Better off/About the same/Worse off. Weighted estimates use sampling probabilities from the realization survey, which are correlated 0.90 and 0.93 with the weight from the paired forecast survey. Sample size is 21,586 in Panel A, 1,679 in Panels B and C, and 1,882 in Panels D and E.

Table E.3: Household financial condition forecast errors are persistent

	Forecast error this survey			
FCE previous survey	Optimist	Realist	Pessimist	Total
Optimist	0.10	0.09	0.01	0.19
Realist	0.08	0.61	0.04	0.73
Pessimist	0.01	0.04	0.03	0.08
Total	0.18	0.74	0.07	1

Note: Sample is 10,546 forecast error pairs from 2,469 panelists. Sample is smaller here than in Appendix Table E.2 because here we require ≥ 2 forecast-realization pairs per panelist and only include realizations of "about the same", to allow for the sharpest feasible test of persistence, by holding realizations constant and allowing for forecast errors in either direction (thereby minimizing measurement error from censoring).

Table E.4: Household financial condition forecast learning?

Panel A. First forecast - realization pair	Realization this year			
Forecast last year	Better	Same	Worse	Total
Better	0.09	0.16	0.06	0.31
Same	0.05	0.40	0.12	0.57
Worse	0.01	0.05	0.06	0.12
Total	0.15	0.61	0.23	1
Panel B. Last forecast - realization pair	Realization this year			
Forecast last year	Better	Same	Worse	Total
Better	0.10	0.13	0.04	0.28
Same	0.06	0.46	0.09	0.61
Worse	0.01	0.05	0.06	0.11
Total	0.17	0.65	0.18	1

Note: Sample includes only the 3073 panelists with multiple forecast-realization pairs.

Table E.5: Pairwise correlations between persistent optimism about financial condition and HtM measures, using all data for non-SZ modules

	Proportion optimistic forecast errors							
	1=($\geq$ 0.5)		1=($>$ 0.5)				Row variable pop. share	
	Unw.	Weighted	Unw.	Weighted	Unw.	Weighted	Unw.	Weighted
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1=(lives paycheck-to-paycheck c. 2012)	0.14	0.21	0.14	0.13	0.14	0.17	0.48	0.50
s.e.	0.05	0.07	0.05	0.07	0.04	0.05	0.02	0.02
N	1068	1068	1068	1068	1068	1068	1068	1068
Lives paycheck-to-paycheck, COVID era	0.19	0.16	0.17	0.11	0.15	0.10	0.38	0.39
s.e.	0.04	0.05	0.04	0.05	0.03	0.03	0.02	0.02
N	1086	1086	1086	1086	1086	1086	1086	1086
1=(Lacks precautionary savings in 2012 and 2018)	0.34	0.32	0.34	0.31	0.30	0.27	0.36	0.39
s.e.	0.05	0.07	0.05	0.07	0.04	0.05	0.02	0.02
N	864	864	864	864	864	864	864	864
1=(Lacks precautionary savings in 2012 or 2018)	0.36	0.34	0.39	0.33	0.36	0.35	0.58	0.62
s.e.	0.05	0.06	0.05	0.07	0.04	0.05	0.02	0.02
N	864	864	864	864	864	864	864	864
Difficulty covering $2k emergency expense	0.17	0.14	0.19	0.15	0.16	0.12	0.48	0.52
s.e.	0.03	0.04	0.03	0.04	0.02	0.03	0.01	0.02
N	2480	2480	2480	2480	2480	2480	2480	2480

Note: Here we combine all the data we have on potentially optimistic financial condition forecast errors and HtM measures. Weighted estimates use the mean sampling weight across all financial condition realizations per panelist.

Table E.6: Pairwise correlations between persistent optimism about financial condition and patience or risk aversion

	Patience		RA: lotteries		RA: scale	
	Unw. (1)	Weighted (2)	Unw. (3)	Weighted (4)	Unw. (5)	Weighted (6)
1=(Prop. optimistic FCEs>0.5)	-0.05	-0.11	-0.05	-0.12	-0.07	-0.20
s.e.	0.07	0.13	0.06	0.10	0.05	0.09
N	447	447	468	468	465	465
1=(Prop. optimistic FCEs$\geq$ 0.5)	-0.01	-0.01	-0.06	-0.12	-0.06	-0.15
s.e.	0.07	0.14	0.06	0.10	0.05	0.09
N	447	447	468	468	465	465
Prop. optimistic forecast errors	-0.12	-0.13	-0.09	-0.15	-0.05	-0.16
s.e.	0.07	0.14	0.06	0.11	0.05	0.08
N	447	447	468	468	465	465

Notes: Persistent optimism measures based on panelists with multiple potentially optimistic forecast errors (see Section 2.1 for details). Patience is the average savings rate across 24 convex time budget choices (Andreoni and Sprenger 2012). Risk aversion (RA) is based on the Barsky et al. (1997) lifetime income gamble elicitation (Columns 3 and 4) or the Dohmen et al. (2010) financial risk-taking scale (Columns 5 and 6). Weighted estimates use sampling probability from the last SZ module. We use Obviously Related Instrumental Variables to account for measurement error in the column variables by using the two measures of each (taken in 2014 and 2017) to instrument for each other (Gillen et al. 2019; Stango and Zinman 2020).

Table E.7: Pairwise correlations between persistent HtM measures and patience or risk aversion

	Patience		RA: lotteries		RA: scale	
	Unw.	Wtd.	Unw.	Wtd.	Unw.	Wtd.
	(1)	(2)	(3)	(4)	(5)	(6)
1=(Persistent severe financial distress)	-0.01	-0.08	0.04	0.08	0.11	0.03
s.e.	(0.06)	(0.14)	(0.05)	(0.12)	(0.04)	(0.09)
N	780	780	832	832	818	818
1=(Persistent low net worth)	-0.03	-0.07	0.14	0.03	0.06	0.08
s.e.	(0.06)	(0.10)	(0.05)	(0.09)	(0.04)	(0.07)
N	734	734	778	778	765	765
1=(paycheck-to-paycheck c. 2012)	0.06	0.38	0.05	-0.16	0.01	0.07
s.e.	(0.10)	(0.17)	(0.09)	(0.31)	(0.07)	(0.16)
N	233	233	260	260	256	256
paycheck-to-paycheck, COVID era	-0.13	-0.01	0.13	0.01	0.08	0.05
s.e.	(0.07)	(0.12)	(0.06)	(0.10)	(0.05)	(0.08
N	493	493	519	519	516	516
1=(Lacks prec. saving in 2012 or 2018)	-0.22	-0.19	0.07	-0.08	0.11	0.05
s.e.	(0.08)	(0.13)	(0.08)	(0.14)	(0.07)	(0.11)
N	254	254	269	269	264	264
Difficult covering \$2k emerg. expenses	-0.15	-0.04	0.11	0.13	0.14	0.15
s.e.	(0.07)	(0.12)	(0.06)	(0.11)	(0.05)	(0.08)
N	462	462	491	491	487	487

Note: Patience is the average savings rate across 24 convex time buget choices (Andreoni and Sprenger 2012). Risk aversion is based on the the the Barsky et al. (1997) lifetime income gamble elicitation (Columns 3 and 4) or the Dohmen et al. (2010) financial risk-taking scale (Columns 5 and 6). Weighted estimates use sampling probability from the last SZ module. We use Obviously Related Instrumental Variables to account for measurement error in the column variables by using the two measures of each (taken in 2014 and 2017) to instrument for each other (Gillen et al. 2019; Stango and Zinman 2020). HtM measures are detailed in Section 5.2. Weighted estimates use the sampling probability for the last SZ module.

Table E.8: Pairwise correlations between overconfidence and patience or risk aversion

	Patience		RA: lotteries		RA: scale	
	Unwtd.	Weighted	Unwtd.	Weighted	Unwtd.	Weighted
	(1)	(2)	(3)	(4)	(5)	(6)
1=Oc both rounds	0.04	-0.01	0.16	0.24	-0.08	-0.20
s.e.	(0.06)	(0.14)	(0.05)	(0.12)	(0.04)	(0.07)
N	758	758	807	807	813	813
Oc percentile rank	0.00	-0.01	0.24	0.31	-0.15	-0.32
s.e.	(0.07)	(0.12)	(0.06)	(0.12)	(0.05)	(0.08)
N	758	758	807	807	813	813

Notes: See Section 5.2 for details on overconfidence measures. Patience is the average savings rate across 24 convex time buget choices Andreoni and Sprenger (2012). Risk aversion is based on the Barsky et al. (1997) lifetime income gamble elicitation (Columns 3 and 4) or the Dohmen et al. (2010) financial risk-taking scale (Columns 5 and 6). Weighted estimates use sampling probability from the last SZ module. We use Obviously Related Instrumental Variables to account for measurement error in the column variables, and in overconfidence percentile rank, by using the two measures of each (taken in 2014 and 2017) to instrument for each other (Gillen et al. 2019; Stango and Zinman 2020).

Table E.9: Persistent overconfidence: Correlations with cognitive skills

	1 = oc both rounds		oc percentile rank	
	Unweighted	Weighted	Unweighted	Weighted
	(1)	(2)	(3)	(4)
Cognitive skill measures				
Summary: 1st common factor	-0.64	-0.63	-0.77	-0.74
s.e.	0.03	0.05	0.04	0.06
N	817	817	817	817
Summary: 1st principal component	-0.55	-0.54	-0.82	-0.83
s.e.	0.03	0.05	0.03	0.05
N	733	733	733	733
Component: Fluid intelligence	-0.72	-0.73	-1.05	-1.07
s.e.	0.03	0.05	0.03	0.06
N	817	817	817	817
Component: Numeracy	-0.36	-0.45	-0.57	-0.66
s.e.	0.04	0.07	0.05	0.08
N	798	798	798	798
Component: Financial literacy	-0.32	-0.24	-0.47	-0.36
s.e.	0.04	0.09	0.04	0.09
N	813	813	813	813
Component: Executive function	-0.32	-0.41	-0.44	-0.60
s.e.	0.05	0.07	0.05	0.09
N	749	749	749	749

Note: Overconfidence re: relative performance in a cognitive skills test (see Section 5.2 for details). All cognitive skills measures are percentile ranks. of each of the component measures shown in the table (see Stango and Zinman (2020) for details on component measures). Weighted estimates use the sampling probability for the last SZ module. All cognitive skills measures, and overconfidence percentile rank, use Obviously Related Instrumental Variables to account for measurement error by having the two rank measures (taken in 2014 and 2017) instrument for each other (Gillen et al. (2019), Stango and Zinman (2020)). We do not take the same approach to the overconfidence indicator in Columns (1) and (2), because measurement error-IV does not work well on misclassification error. An IV point estimate of a correlation can exceed |1|, as it does in two instances here.

Figure E.11: Distribution of HtM along the income distribution

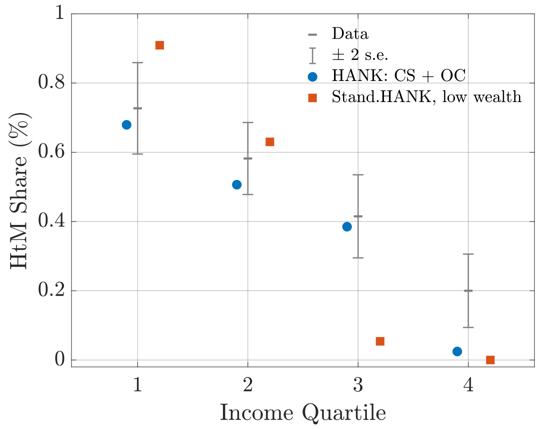

Note: This figure shows the share of hand-to-mouth households along the income distribution for our "low net worth" HtM measure in the data (black dashes). It also shows the share of HtM households in our baseline model with overconfidence (blue dots) and in the standard HANK model recalibrated to match the average MPC of our baseline model (red squares). We redefine the HtM measure in both models such that the aggregate HtM share matches the data.

E.2 Proofs

[Proof of Lemma 3] Lemma 3 says that unless marginal utility is constant across income states, heterogeneity in overconfidence and heterogeneity in patience are not equivalent. To see this, consider a simple counterexample. Focus on two households, $i \in \{1, 2\}$, and two possible future states, which we denote by U and D (e.g., for Up and Down). We focus on the equivalence of overconfident households and relatively impatient households with a discount factor $\hat{\beta} < \beta$. If overconfidence and patience heterogeneity are equivalent, it has to hold that the Euler equations of unconstrained households have to be identical. Imposing that household 1 has the same marginal utility in both economies in the current period implies that the expected discounted future marginal utility has to be identical, too:

$$\beta \tilde{E}_t \left[u'(c^1_{t+1}) \right] = \hat{\beta} E_t \left[u'(\hat{c}^1_{t+1}) \right], \tag{E.1}$$

where a hat "$\hat{\cdot}$" denotes the economy with heterogeneity in patience. Similarly, for household 2:

$$\beta \tilde{E}_t \left[u'(c^2_{t+1}) \right] = \hat{\beta} E_t \left[u'(\hat{c}^2_{t+1}) \right], \tag{E.2}$$

Assuming, without loss of generality, that household 1 starts in the U state and denoting the probability of moving to the D state by p_{UD}, equation (E.1) implies

$$\frac{\beta}{\hat{\beta}} = \frac{p_{UD} u'(c^{1,D}_{t+1}) + (1 - p_{UD}) u'(c^{1,U}_{t+1})}{\frac{1}{\alpha} p_{UD} u'(c^{1,D}_{t+1}) + (1 - \frac{1}{\alpha} p_{UD}) u'(c^{1,U}_{t+1})}. \tag{E.3}$$

(Implicitly, but without loss of generality, we assume here that consumption in the U state is higher than in the D state). Similarly, for household 2, who starts in state D

$$\frac{\beta}{\hat{\beta}} = \frac{p_{DU} u'(c^{2,U}_{t+1}) + (1 - p_{DU}) u'(c^{2,D}_{t+1})}{\alpha p_{DU} u'(c^{2,U}_{t+1}) + (1 - \alpha p_{DU}) u'(c^{2,D}_{t+1})}. \tag{E.4}$$

Thus, for given transition probabilities, degree of overconfidence α, discount factor in the economy with overconfidence β, and marginal utilities across states, we have one free parameter, $\hat{\beta}$, but two equations that need to hold.[109] Thus, the two economies are in general not identical (it becomes even less likely that the two economies are identical when we allow for more states and households). The only case in which the two are identical is when marginal utility is constant across states, that is when households can perfectly insure themselves against income shocks. Given our incomplete-markets setup, however, that is generally not the case, and therefore, heterogeneity in overconfidence is not equivalent to heterogeneity in patience.

[109] A simple numerical example illustrates this. Assume $p_{UD} = p_{DU} = 0.5$, $\alpha = 2$, $u'(c^D) = 1$ and $u'(c^U) = 2 > 1$. It follows that equation (E.1) implies a discount factor ratio of 0.86 whereas equation (E.1) implies a ratio of 0.75.

Bibliography

Acharya, S., E. Challe, and K. Dogra (2023). Optimal monetary policy according to hank. *American Economic Review 113*(7), 1741–1782.

Acharya, S. and K. Dogra (2020). Understanding hank: Insights from a prank. *Econometrica 88*(3), 1113–1158.

Adam, K., O. Pfaeuti, and T. Reinelt (2022). Subjective housing price expectations, falling natural rates and the optimal inflation target.

Afrouzi, H., S. Y. Kwon, A. Landier, Y. Ma, and D. Thesmar (2022). Overreaction in expectations: Evidence and theory. *The Quarterly Journal of Economics (forthcoming)*.

Agarwal, S. and B. Mazumder (2013). Cognitive abilities and household financial decision making. *American Economic Journal: Applied Economics 5*(1), 193–207.

Aggarwal, R., A. Auclert, M. Rognlie, and L. Straub (2023). Excess savings and twin deficits: The transmission of fiscal stimulus in open economies. *NBER Macroeconomics Annual 37*(1), 325–412.

Aguiar, M. A., M. Bils, and C. Boar (forthcoming). Who are the hand-to-mouth?. *Review of Economic Studies*.

Ahn, S., G. Kaplan, B. Moll, T. Winberry, and C. Wolf (2018). When inequality matters for macro and macro matters for inequality. *NBER macroeconomics annual 32*(1), 1–75.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics 109*(3), 659–684.

Aiyagari, S. R. and E. R. McGrattan (1998). The optimum quantity of debt. *Journal of Monetary Economics 42*(3), 447–469.

Alesina, A. and R. J. Barro (2002, 05). Currency Unions. *The Quarterly Journal of Economics 117*(2), 409–436.

Almgren, M., J.-E. Gallegos, J. Kramer, and R. Lima (2022). Monetary policy and liquidity constraints: Evidence from the euro area. *American Economic Journal: Macroeconomics 14*(4), 309–340.

Ampudia, M., D. Georgarakos, J. Slacalek, O. Tristani, P. Vermeulen, and G. Violante (2018). Monetary policy and household inequality.

Andrade, P., G. Gaballo, E. Mengus, and B. Mojon (2019). Forward guidance and heterogeneous beliefs. *American Economic Journal: Macroeconomics 11*(3), 1–29.

Andreoni, J. and C. Sprenger (2012). Estimating time preferences from convex budgets.

American Economic Review 102(7), 3333–3356.

Andreou, E., M. Demetriadou, and A. Tryphonides (2023). The cross section of household preferences and the marginal propensity to consume: Evidence from high-frequency data. *Working Paper*.

Angeletos, G.-M., F. Collard, and H. Dellas (2020). Business-cycle anatomy. *American Economic Review 110*(10), 3030–70.

Angeletos, G.-M., F. Collard, and H. Dellas (2023). Public debt as private liquidity: Optimal policy. *Journal of Political Economy 131*(11), 3233–3264.

Angeletos, G.-M. and Z. Huo (2021). Myopia and anchoring. *American Economic Review 111*(4), 1166–1200.

Angeletos, G.-M., Z. Huo, and K. A. Sastry (2021). Imperfect macroeconomic expectations: Evidence and theory. *NBER Macroeconomics Annual 35*(1), 1–86.

Angeletos, G.-M. and C. Lian (2017). Dampening general equilibrium: From micro to macro. Technical report, National Bureau of Economic Research.

Angeletos, G.-M. and C. Lian (2018). Forward guidance without common knowledge. *American Economic Review 108*(9), 2477–2512.

Angeletos, G.-M. and C. Lian (2023). Dampening general equilibrium: incomplete information and bounded rationality. In *Handbook of Economic Expectations*, pp. 613–645. Elsevier.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review 109*(6), 2333–67.

Auclert, A., B. Bardóczy, and M. Rognlie (2021). Mpcs, mpes, and multipliers: A trilemma for new keynesian models. *The Review of Economics and Statistics*, 1–41.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica 89*(5), 2375–2408.

Auclert, A. and M. Rognlie (2020). Inequality and aggregate demand.

Auclert, A., M. Rognlie, M. Souchier, and L. Straub (2021). Exchange rates and monetary policy with heterogeneous agents: Sizing up the real income channel. mimeo.

Auclert, A., M. Rognlie, and L. Straub (2018, September). The Intertemporal Keynesian Cross. NBER Working Paper No. 25020.

Auclert, A., M. Rognlie, and L. Straub (2020). Micro jumps, macro humps: Monetary policy and business cycles in an estimated hank model. Technical report, National Bureau of Economic Research.

Auclert, A., M. Rognlie, and L. Straub (forthcoming). The intertemporal keynesian cross. *Journal of Political Economy*.

Bachmann, R., D. Baqaee, C. Bayer, M. Kuhn, A. Löschel, B. Moll, A. Peichl, K. Pittel, and M. Schularick (2022). What if? the economic effects for germany of a stop of energy imports from russia. ECONtribute Policy Brief No. 028.

Bachmann, R., B. Born, O. Goldfayn-Frank, G. Kocharkov, R. Luetticke, and M. Weber (2021). A temporary vat cut as unconventional fiscal policy. NBER Working Paper 29442.

Baker, S. R., L. Kueng, L. McGranahan, and B. T. Melzer (2019). Do household finances constrain unconventional fiscal policy? *Tax Policy and the Economy 33*(1), 1–32.

Balleer, A., G. Duernecker, S. Forstner, and J. Goensch (2022). The effects of biased labor market expectations on consumption, wealth inequality, and welfare. *Working Paper*.

Banks, J. and Z. Oldfield (2007). Understanding pensions: Cognitive function, numerical ability and retirement saving. *Fiscal Studies 28*(2), 143–170.

Barsky, R. B., F. T. Juster, M. S. Kimball, and M. D. Shapiro (1997). Preference parameters and behavioral heterogeneity: An experimental approach in the health and retirement study. *The Quarterly Journal of Economics 112*(2), 537–579.

Bayer, C., B. Born, and R. Luetticke (2020a). The liquidity channel of fiscal policy. *CEPR Discussion Paper No. DP14883*.

Bayer, C., B. Born, and R. Luetticke (2020b). Shocks, frictions, and inequality in US business cycles. CEPR Discussion Paper No. 14364.

Bayer, C., B. Born, and R. Luetticke (2023). The liquidity channel of fiscal policy. *Journal of Monetary Economics 134*, 86–117.

Bayer, C., B. Born, and R. Luetticke (forthcoming). Shocks, frictions, and inequality in us business cycles.

Bayer, C., B. Born, R. Luetticke, and G. J. Müller (2020). The coronavirus stimulus package: How large is the transfer multiplier?

Bayer, C., B. Born, R. Luetticke, and G. J. Müller (2023). The coronavirus stimulus package: How large is the transfer multiplier. *The Economic Journal 133*(652), 1318–1347.

Bayer, C., A. Kriwoluzky, G. Müller, and F. Seyrich (2023, July). A HANK2 model of monetary unions. CEPR Discussion Paper No. 18258.

Bayer, C., A. Kriwoluzky, G. J. Müller, and F. Seyrich (2022). Das Sein bestimmt das Bewußtsein: Attitudes towards fiscal policy according to $HANK^2$. mimeo.

Bayer, C., A. Kriwoluzky, G. J. Müller, and F. Seyrich (2023). Hicks in hank: Fiscal responses to an energy shock.

Bayer, C. and R. Luetticke (2020). Solving discrete time heterogeneous agent models with aggregate risk and many idiosyncratic states by perturbation. *Quantitative Economics 11*(4), 1253–1288.

Bayer, C., R. Lütticke, L. Pham-Dao, and V. Tjaden (2019). Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica 87*(1), 255–290.

Bayoumi, T. and B. Eichengreen (1992, January). Shocking Aspects of European Monetary Unification. NBER Working Papers 3949.

Beetsma, R. and H. Jensen (2005). Monetary and fiscal policy interactions in a micro-founded model of a monetary union. *Journal of International Economics 67*(2), 320–352.

Beetsma, R. and H. Uhlig (1999). An analysis of the stability and growth pact. *Economic Journal 109*(458), 546–71.

Bellifemine, M., A. Couturier, and R. Jamilov (2023). The Regional Keynesian Cross. mimeo.

Benigno, P. (2004). Optimal monetary policy in a currency area. *Journal of International Economics 63*(2), 293–320.

Benigno, P. and J. D. López-Salido (2006). Inflation persistence and optimal monetary policy in the euro area. *Journal of Money, Credit and Banking 38*(3), 587–614.

Benigno, P. and L. Paciello (2014). Monetary policy, doubts and asset prices. *Journal of Monetary Economics 64*, 85–98.

Beraja, M. and N. Zorzi (2024). Durables and size-dependence in the marginal propensity to spend. *Working Paper*.

Bernard, J.-T. and M. Kichian (2021). The impact of a revenue-neutral carbon tax on gdp dynamics: the case of british columbia. *The Energy Journal 42*(3), 205–224.

Bernard, R. (2023). Mental accounting and the marginal propensity to consume. *Working Paper*.

Bewley, T. (1986). Stationary monetary equilibrium with a continuum of independently fluctuating consumers. *Contributions to mathematical economics in honor of Gérard Debreu 79*.

Bhandari, A., J. Borovička, and P. Ho (2019). Survey data and subjective beliefs in business cycle models.

Bhandari, A., D. Evans, M. Golosov, and T. J. Sargent (2021). Inequality, business cycles, and monetary-fiscal policy. *Econometrica 89*(6), 2559–2599.

Bianchi-Vimercati, R., M. S. Eichenbaum, and J. Guerreiro (2021). Fiscal policy at the zero lower bound without rational expectations. *NBER Working Paper Series No. 29134*.

Bilbiie, F. O. (2008). Limited asset markets participation, monetary policy and (inverted) aggregate demand logic. *Journal of economic theory 140*(1), 162–196.

Bilbiie, F. O. (2020). The new keynesian cross. *Journal of Monetary Economics 114*, 90–108.

Bilbiie, F. O. (2021). Monetary policy and heterogeneity: An analytical framework.

Bilbiie, F. O., D. R. Känzig, and P. Surico (2022). Capital and income inequality: An aggregate-demand complementarity. *Journal of Monetary Economics 126*, 154–169.

Boppart, T., P. Krusell, and K. Mitman (2018). Exploiting mit shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control 89*, 68–92.

Bordalo, P., N. Gennaioli, Y. Ma, and A. Shleifer (2020). Overreaction in macroeconomic expectations. *American Economic Review 110*(9), 2748–82.

Born, B., Z. Enders, M. Menkhoff, G. J. Müller, and K. Niemann (2022). Firm expectations and news: Micro v macro.

Boutros, M. (2022). Windfall income shocks with finite planning horizons.

Broer, T., N.-J. Harbo Hansen, P. Krusell, and E. Öberg (2020). The new keynesian transmission mechanism: A heterogeneous-agent perspective. *The Review of Economic Studies 87*(1), 77–101.

Broer, T., A. Kohlhas, K. Mitman, and K. Schlafmann (2022). Information and wealth heterogeneity in the macroeconomy.

Brunnermeier, M. K. and J. A. Parker (2005). Optimal expectations. *American Economic Review 95*(4), 1092–1118.

Caballero, R. J. and A. Simsek (2019). Prudential monetary policy. Technical report, National Bureau of Economic Research.

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of monetary Economics 12*(3), 383–398.

Caplin, A., V. Gregory, E. Lee, S. Leth-Petersen, and J. Sæverud (2024). Subjective earnings risk. *Working Paper*.

Carlson, J. A. and M. Parkin (1975). Inflation expectations. *Economica 42*(166), 123–138.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics letters 91*(3), 312–320.

Castaneda, A., J. Dıaz-Giménez, and J.-V. Rıos-Rull (1998). Exploring the income distribution business cycle dynamics. *Journal of Monetary economics 42*(1), 93–130.

Chapman, J., M. Dean, P. Ortoleva, E. Snowberg, and C. Camerer (2023). Econographics. *Journal of Political Economy Microeconomics 1*(1), 115–161.

Chen, X., S. Lazarakis, and P. Varthalitis (2023). Debt targets and fiscal consolidation in a two-country HANK model for Eurozone. mimeo.

Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica 80*(3), 969–1018.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011). Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. *American Economic Review 101*(3), 471–475.

Christiano, L., M. Eichenbaum, and S. Rebelo (2011). When is the government spending multiplier large? *Journal of Political Economy 119*(1), 78–121.

Claus, E. and V. H. Nguyen (2023). Biased expectations. *European Economic Review 154*, 104424.

Cochrane, J. H. (2018). Michelson-morley, fisher, and occam: The radical implications of stable quiet inflation at the zero bound. *NBER Macroeconomics Annual 32*(1), 113–226.

Coibion, O. and Y. Gorodnichenko (2015). Information rigidity and the expectations formation process: A simple framework and new facts. *American Economic Review 105*(8), 2644–78.

Coibion, O., Y. Gorodnichenko, E. S. Knotek II, and R. Schoenle (2022). Average inflation targeting and household expectations. *Journal of Political Economy Macroeconomics (forthcoming).*

Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia (2017). Innocent bystanders? monetary policy and inequality. *Journal of Monetary Economics 88*, 70–89.

Correia, I., E. Farhi, J. P. Nicolini, and P. Teles (2013). Unconventional fiscal policy at the zero bound. *American Economic Review 103*(4), 1172–1211.

Correia, I., J. P. Nicolini, and P. Teles (2008). Optimal fiscal and monetary policy: Equivalence results. *Journal of Political Economy 116*(1), 141–170.

Corsetti, G. (2008, March). A Modern Reconsideration of the Theory of Optimal Currency Areas. Technical report.

Corsetti, G., A. Meier, and G. J. Müller (2012). Fiscal Stimulus with Spending Reversals. *The Review of Economics and Statistics 94*(4), 878–895.

Dalton, P. S., S. Ghosal, and A. Mani (2016). Poverty and aspirations failure. *The Economic Journal 126*(590), 165–188.

Dávila, E. and A. Schaab (2022). Optimal monetary policy with heterogeneous agents: A timeless ramsey approach. *Available at SSRN 4102028*.

Dávila, E. and A. Schaab (2023). Welfare assessments with heterogeneous individuals. *Working Paper*.

Davila, J., J. H. Hong, P. Krusell, and J.-V. Ríos-Rull (2012). Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks. *Econometrica 80*(6), 2431–2467.

de Ferra, S., K. Mitman, and F. Romei (2021). Why does capital flow from equal to unequal countries? *CEPR Discussion paper 15647*.

Debortoli, D. and J. Galí (2018). Monetary policy with heterogeneous agents: Insights from tank models.

Debortoli, D., J. Galí, and L. Gambetti (2020). On the empirical (ir) relevance of the zero lower bound constraint. *NBER Macroeconomics Annual 34*(1), 141–170.

Del Negro, M., M. Giannoni, and C. Patterson (2015). The forward guidance puzzle. *Journal of Political Economy Macroeconomics (forthcoming)*.

Dengler, T. and B. Gehrke (2023). Short-time work and precautionary savings. *Working Paper*.

d'Haultfoeuille, X., C. Gaillac, and A. Maurel (2021). Rationalizing rational expectations: Characterizations and tests. *Quantitative Economics 12*(3), 817–842.

di Giovanni, J., Kalemli-Özcan, A. Silva, and M. A. Yildirim (2022). Global supply chain pressures, international trade, and inflation.

Dohmen, T., A. Falk, D. Huffman, and U. Sunde (2010). Are risk aversion and impatience related to cognitive ability? *American Economic Review 100*(3), 1238–1260.

D'Acunto, F., D. Hoang, M. Paloviita, and M. Weber (2019a). Cognitive abilities and inflation expectations. *AEA Papers and Proceedings 109*, 562–66.

D'Acunto, F., D. Hoang, M. Paloviita, and M. Weber (2019b). Iq, expectations, and choice. Technical report, National Bureau of Economic Research.

D'Acunto, F., D. Hoang, M. Paloviita, and M. Weber (2023). Cognitive constraints and economic incentives. *Working Paper*.

D'Acunto, F., D. Hoang, and M. Weber (2022). Managing households' expectations with unconventional policies. *The Review of Financial Studies 35*(4), 1597–1642.

Ehrlinger, J., K. Johnson, M. Banner, D. Dunning, and J. Kruger (2008). Why the unskilled are unaware: Further explorations of (absent) self-insight among the incompetent. *Organizational Behavior and Human Decision Processes 105*(1), 98–121.

Enders, Z., P. Jung, and G. J. Müller (2013). Has the euro changed the business cycle? *European Economic Review 59*, 189–211.

Erceg, C. J., D. W. Henderson, and A. T. Levin (2000). Optimal monetary policy with staggered wage and price contracts. *Journal of monetary Economics 46*(2), 281–313.

Farhi, E. and I. Werning (2017). Fiscal unions. *American Economic Review 107*(12), 3788–3834.

Farhi, E. and I. Werning (2019). Monetary policy, bounded rationality, and incomplete markets. *American Economic Review 109*(11), 3887–3928.

Feldstein, M. (2002). The role for discretionary fiscal policy in a low interest rate environment. *NBER Working Paper Series No. 9203*.

Ferriere, A. and G. Navarro (2022). The heterogeneous effects of government spending: It's all about taxes.

Floden, M. and J. Lindé (2001). Idiosyncratic risk in the united states and sweden: Is there a role for government insurance? *Review of Economic dynamics 4*(2), 406–437.

Frankel, J. A. and A. K. Rose (1998). The endogeneity of the optimum currency area criteria. *The Economic Journal 108*(449), 1009–1025.

Fuhrer, J. C. and G. D. Rudebusch (2004). Estimating the euler equation for output. *Journal of Monetary Economics 51*(6), 1133–1153.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics 129*(4), 1661–1710.

Gabaix, X. (2017). Behavioral macroeconomics via sparse dynamic programming.

Gabaix, X. (2019). Behavioral inattention. In *Handbook of Behavioral Economics: Applications and Foundations 1*, Volume 2, pp. 261–343. Elsevier.

Gabaix, X. (2020). A behavioral new keynesian model. *American Economic Review 110*(8), 2271–2327.

Galí, J. (2015). *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications.* Princeton University Press.

Galí, J. and T. Monacelli (2005). Monetary policy and exchange rate volatility in a small open economy. *Review of Economic Studies 72*, 707–734.

Gali, J. and T. Monacelli (2008). Optimal monetary and fiscal policy in a currency union. *Journal of International Economics 76*(1), 116–132.

Galí, J. and T. Monacelli (2016). Understanding the gains from wage flexibility: The exchange rate connection. *American Economic Review 106*(12), 3829–68.

Genicot, G. and D. Ray (2017). Aspirations and inequality. *Econometrica 85*(2), 489–519.

Gillen, B., E. Snowberg, and L. Yariv (2019). Experimenting with measurement error: Techniques with applications to the caltech cohort study. *Journal of Political Economy 127*(4), 1826–1863.

Goulder, L. and M. Hafstead (2017). *Confronting the climate challenge: US policy options.* Columbia University Press.

Goulder, L. H. (1995). Environmental taxation and the double dividend: a reader's guide. *International tax and public finance 2*, 157–183.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988). Investment, capacity utilization, and the real business cycle. *The American Economic Review*, 402–417.

Grohmann, A., L. Menkhoff, C. Merkle, and R. Schmacker (2023). Earn more tomorrow: Overconfidence, income expectations and consumer indebtedness. *Working Paper*.

Groll, D. and T. Monacelli (2020). The inherent benefit of monetary unions. *Journal of Monetary Economics 111*(C), 63–79.

Guerreiro, J. (2023). Belief disagreement and business cycles. *Working Paper*.

Guerrieri, V. and G. Lorenzoni (2017). Credit crises, precautionary savings, and the liquidity trap. *The Quarterly Journal of Economics 132*(3), 1427–1467.

Guo, X., P. Ottonello, and D. J. Perez (2020). Monetary policy and redistribution in open economies. mimeo, National Bureau of Economic Research.

Guvenen, F., S. Ozkan, and J. Song (2014). The nature of countercyclical income risk. *Journal of Political Economy 122*(3), 621–660.

Hagedorn, M. (2016). *A demand theory of the price level.*

Hagedorn, M. (2018). Prices and inflation when government bonds are net wealth.

Hagedorn, M., J. Luo, I. Manovskii, and K. Mitman (2019). Forward guidance. *Journal of Monetary Economics 102*, 1–23.

Hagedorn, M., I. Manovskii, and K. Mitman (2019). The fiscal multiplier. *National Bureau of Economic Research.*

Hall, R. E. (2011). The long slump. *American Economic Review 101*(2), 431–69.

Hanushek, E. A. and L. Woessmann (2008). The role of cognitive skills in economic development. *Journal of Economic Literature 46*(3), 607–668.

Havranek, T. and A. Sokolova (2020). Do consumers really follow a rule of thumb? three thousand estimates from 144 studies say "probably not". *Review of Economic Dynamics 35*, 97–122.

Heidhues, P., B. Kőszegi, and P. Strack (2018). Unrealistic expectations and misguided learning. *Econometrica 86*(4), 1159–1214.

Hettig, T. and G. J. Müller (2018). Fiscal policy coordination in currency unions at the effective lower bound. *Journal of International Economics 115*(C), 80–98.

Hintermaier, T. and W. Koeniger (2019). Differences in euro-area household finances and their relevance for monetary-policy transmission.

Hoffman, M. and S. V. Burks (2020). Worker overconfidence: Field evidence and implications for employee turnover and firm profits. *Quantitative Economics 11*(1), 315–348.

Holm, M. B., P. Paul, and A. Tischbirek (2021). The transmission of monetary policy under the microscope. *Journal of Political Economy 129*(10), 2861–2904.

Huffman, D., C. Raymond, and J. Shvets (2022). Persistent overconfidence and biased memory: Evidence from managers. *American Economic Review 112*(10), 3141–75.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of economic Dynamics and Control 17*(5), 953–969.

Hyytinen, A. and H. Putkuri (2018). Household optimism and overborrowing. *Journal of Money, Credit and Banking 50*(1), 55–76.

Ilut, C. and R. Valchev (2023). Economic agents as imperfect problem solvers. *The Quarterly Journal of Economics 138*(1), 313–362.

Jappelli, T. and L. Pistaferri (2010). The consumption response to income changes. *Annu. Rev. Econ. 2*(1), 479–506.

Jensen, A. R. (1998). The g factor: The science of mental ability. *Westport, CT: Prager*.

Jordà, Ò., K. Knoll, D. Kuvshinov, M. Schularick, and A. M. Taylor (2019). The rate of return on everything, 1870–2015. *The Quarterly Journal of Economics 134*(3), 1225–1298.

Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to hank. *American Economic Review 108*, 697–743.

Kaplan, G. and G. L. Violante (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica 82*(4), 1199–1239.

Kaplan, G. and G. L. Violante (2022). The marginal propensity to consume in heterogeneous agent models. *Annual Review of Economics 14*, 747–775.

Kekre, R. (2021). Unemployment insurance in macroeconomic stabilization. *NBER Working Paper Series No. 29505*.

Kekre, R. and M. Lenel (2022). Monetary policy, redistribution, and risk premia. *Econometrica 90*(5), 2249–2282.

Kenen, P. B. (1981). *The Theory of Optimum Currency Areas: An Eclectic View*, pp. 163–182. Princeton: Princeton University Press.

Kindermann, F. and S. Kohls (2017). Rental markets and wealth inequality across europe. mimeo.

Koby, Y. and C. Wolf (2020). Aggregation in heterogeneous-firm models: Theory and measurement. *Manuscript, July*.

Kollmann, R. (2004). Welfare effects of a monetary union: The role of trade openness. *Journal of the European Economic Association 2*(2/3), 289–301.

Konradt, M. and B. Weder (2021). Carbon taxation and inflation: Evidence from the european and canadian experience. Technical report, Graduate Institute of International and Development Studies Working Paper.

Krueger, D., K. Mitman, and F. Perri (2016). Macroeconomics and household heterogeneity. In *Handbook of Macroeconomics*, Volume 2, pp. 843–921. Elsevier.

Krugman, P. (1993). Lessons of massachusetts for emu. In F. Giavazzi and F. Torres (Eds.), *Adjustment and growth in the European Monetary Union*, Chapter 8, pp. 241–261. Cambridge University Press.

Kučinskas, S. and F. S. Peters (2022). Measuring under-and overreaction in expectation formation. *The Review of Economics and Statistics*, 1–45.

Kuhn, M. and L. Schlattmann (2024). Distributional consequences of climate policies. *Distributional Consequences of Climate Policies*.

Kuhn, M., M. Schularick, and U. I. Steins (2020). Income and wealth inequality in america, 1949–2016. *Journal of Political Economy 128*(9), 3469–3519.

Känzig, D. R. (2023, May). The unequal economic consequences of carbon pricing. NBER Working Paper No. 31221.

Laibson, D., P. Maxted, and B. Moll (2021). Present bias amplifies the household balance-sheet channels of macroeconomic policy.

Langot, F., S. Malmberg, F. Tripier, and J.-O. Hairault (2023). Making environmental policies acceptable. Technical report, Mimeo, CEPREMAP.

Le Grand, F., A. Martin-Baillon, and X. Ragot (2021). Should monetary policy care about redistribution? optimal fiscal and monetary policy with heterogeneous agents.

Lian, C. (2021). Mistakes in future consumption, high mpcs now. Technical report, National Bureau of Economic Research.

Luetticke, R. (2021). Transmission of monetary policy with heterogeneity in household portfolios. *American Economic Journal: Macroeconomics 13*(2), 1–25.

Lusardi, A. and O. S. Mitchell (2014). The economic importance of financial literacy: Theory and evidence. *American Economic Journal: Journal of Economic Literature 52*(1), 5–44.

MacLeod, C. M. (1991). Half a century of research on the stroop effect: An integrative review. *Psychological Bulletin 109*(2), 163.

Mankiw, N. G. (2000). The savers-spenders theory of fiscal policy. *American economic review 90*(2), 120–125.

McArdle, J. J., G. G. Fisher, and K. M. Kadlec (2007). Latent variable analyses of age trends of cognition in the health and retirement study, 1992-2004. *Psychology and Aging 22*(3), 525.

McClung, N. and T. Nighswander (2021). Belief heterogeneity and the income-wealth relationship. *Working Paper*.

McKay, A., E. Nakamura, and J. Steinsson (2016). The power of forward guidance revisited. *American Economic Review 106*(10), 3133–58.

McKay, A., E. Nakamura, and J. Steinsson (2017). The discounted euler equation: A note. *Economica 84*(336), 820–831.

McKay, A. and J. F. Wieland (2022). Forward guidance and durable goods demand.

American Economic Review: Insights 4(1), 106–22.

McKay, A. and C. Wolf (2022). Optimal policy rules in hank. mimeo.

McKibbin, W. J., A. C. Morris, A. Panton, and P. Wilcoxen (2017). Climate change and monetary policy: Dealing with disruption.

McKinnon, R. I. (1963). Optimum currency areas. *The American Economic Review 53*(4), 717–725.

Metcalf, G. E. (2019). On the economics of a carbon tax for the united states. *Brookings Papers on Economic Activity 2019*(1), 405–484.

Metcalf, G. E. and J. H. Stock (2023). The macroeconomic impact of europe's carbon taxes. *American Economic Journal: Macroeconomics 15*(3), 265–286.

Michaillat, P. and E. Saez (2021). Resolving new keynesian anomalies with wealth in the utility function. *Review of Economics and Statistics 103*(2), 197–215.

Miyake, A. and N. P. Friedman (2012). The nature and organization of individual differences in executive functions: Four general conclusions. *Current Directions in Psychological Science 21*(1), 8–14.

Moschini, E., G. Raveendranathan, and M. Xu (2023). Optimism about graduation and college financial aid. *Working Paper*.

Mueller, A. I., J. Spinnewijn, and G. Topa (2021). Job seekers' perceptions and employment prospects: Heterogeneity, duration dependence, and bias. *American Economic Review 111*(1), 324–363.

Mumtaz, H. and A. Theophilopoulou (2017). The impact of monetary policy on inequality in the uk. an empirical analysis. *European Economic Review 98*, 410–423.

Mundell, R. A. (1961). A theory of optimum currency areas. *The American Economic Review 51*(4), 657–665.

Oh, H. and R. Reis (2012). Targeted transfers and the fiscal response to the great recession. *Journal of Monetary Economics 59*, 50–64.

Oskolkov, A. (2023). Exchange rate policy and heterogeneity in small open economies. *Journal of International Economics 142*, 103750.

Ottonello, P. and T. Winberry (2020). Financial heterogeneity and the investment channel of monetary policy. *Econometrica 88*(6), 2473–2502.

Patterson, C. (2023). The matching multiplier and the amplification of recessions. *American Economic Review 113*(4), 982–1012.

Pfäuti, O. and F. Seyrich (2023). A behavioral heterogeneous agent new keynesian model. *Working Paper*.

Pfäuti, O., F. Seyrich, and J. Zinman (2024). Bad luck or bad decisions? macroeconomic implications of persistent heterogeneity in cognitive skills and overconfidence. Technical report, National Bureau of Economic Research.

Pham-Dào, L. A. (2016). *Income Risk, Self-Insurance and Public Policies: Consequences for Aggregate Fluctuations and Wealth Inequality.* Ph. D. thesis, Universitäts-und Landesbibliothek Bonn.

Ravn, M. O. and V. Sterk (2021). Macroeconomic fluctuations with hank & sam: An analytical approach. *Journal of the European Economic Association 19*(2), 1162–1202.

Reiter, M., T. Sveen, and L. Weinke (2013). Lumpy investment and the monetary transmission mechanism. *Journal of Monetary Economics 60*(7), 821–834.

Rose, A. K. (2000). One money, one market: the effect of common currencies on trade. *Economic Policy 15*(30), 08–45.

Roth, C., M. Wiederholt, and J. Wohlfart (2021). The effects of forward guidance: Theory with measured expectations.

Rouwenhorst, K. G. (1995). Asset pricing implications of equilibrium business cycle models. pp. 294–330.

Rozsypal, F. and K. Schlafmann (forthcoming). Overpersistence bias in individual income expectations and its aggregate implications. *American Economic Journal: Macroeconomics*.

Samarina, A. and A. D. Nguyen (2019). Does monetary policy affect income inequality in the euro area?

Seidl, H. and F. Seyrich (2023). Unconventional fiscal policy in a heterogeneous-agent new keynesian model. *Journal of Political Economy Macroeconomics*, forthcoming.

Sergeyev, D., C. Lian, and Y. Gorodnichenko (2024). The economics of financial stress. *Working Paper*.

Slacalek, J., O. Tristani, and G. L. Violante (2020). Household balance sheet channels of monetary policy: A back of the envelope calculation for the euro area. *Journal of Economic Dynamics and Control 115*, 103879.

Smirnov, D. (2022). Optimal monetary policy in hank. *Working paper*.

Solon, G., S. J. Haider, and J. M. Wooldridge (2015). What are we weighting for? *Journal of Human Resources 50*(2), 301–316.

Souleles, N. S. (2004). Expectations, heterogeneous forecast errors, and consumption: Micro evidence from the michigan consumer sentiment surveys. *Journal of Money, Credit and Banking*, 39–72, Vol. 36(1).

Stango, V. and J. Zinman (2020). We are all behavioral, more or less: A taxonomy of consumer decision making. Technical report, National Bureau of Economic Research.

Stango, V. and J. Zinman (2024). Behavioral biases are temporally stable. *Working Paper*.

Storesletten, K., C. I. Telmer, and A. Yaron (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of political Economy 112*(3), 695–717.

Taylor, J. B. (1993). Discretion versus policy rules in practice. In *Carnegie-Rochester conference series on public policy*, Volume 39, pp. 195–214. Elsevier.

Wang, T. (2023). Perceived versus calibrated income risks in heterogeneous-agent consumption models. *Working Paper*.

Werning, I. (2015). Incomplete markets and aggregate demand.

Wiederholt, M. (2015). Empirical properties of inflation expectations and the zero lower bound. *manuscript, Goethe University*.

Wolf, C. (2021). Interest rate cuts vs. stimulus payments: An equivalence result.

Woodford, M. (1990). Public debt as private liquidity. *The American Economic Review 80*(2), 382–388.

Woodford, M. (2003). Interest and prices.

Woodford, M. (2013). Macroeconomic analysis without the rational expectations hypothesis. *Annual Review of Economics 5*(1), 303–346.

Woodford, M. (2019). Monetary policy analysis when planning horizons are finite. *NBER Macroeconomics Annual 33*(1), 1–50.

Yang, Y. (2022). Redistributive inflation and optimal monetary policy. *Working paper*.

Young, E. R. (2010). Solving the incomplete markets model with aggregate uncertainty using the krusell–smith algorithm and non-stochastic simulations. *Journal of Economic Dynamics and Control 34*(1), 36–41.

Yun, T. (1996). Nominal price rigidity, money supply endogeneity, and business cycles. *Journal of monetary Economics 37*(2), 345–370.

Zhou, H. (2021). Open economy, redistribution, and the aggregate impact of external shocks. mimeo.

Zimmermann, F. (2020). The dynamics of motivated beliefs. *American Economic Review 110*(2), 337–363.

Erklärungen

Erklärung gem. §4 PO

Hiermit erkläre ich, dass ich mich noch keinem Promotionsverfahren unterzogen oder um Zulassung zu einem solchen beworben habe, und die Dissertation in der gleichen oder einer anderen Fassung bzw. Überarbeitung einer anderen Fakultät, einem Prüfungsausschuss oder einem Fachvertreter an einer anderen Hochschule nicht bereits zur Überprüfung vorgelegen hat.

Berlin, den 11. Juni 2024

. . . .

Fabian Georg Ludwig Seyrich

Erklärung gem. §10 Abs. 3 PO

Hiermit erkläre ich, dass ich für die Dissertation folgende Hilfsmittel und Hilfen verwendet habe:

- MATLAB
- Julia
- Excel
- LaTeX
- DeepL
- Siehe References, für die verwendete Literatur

Auf dieser Grundlage habe ich die Arbeit selbstständig verfasst.

Berlin, den 11. Juni 2024

. . . .

Fabian Georg Ludwig Seyrich